

Model instability in predictive exchange rate regressions

Hauzenberger, Niko; Huber, Florian

DOI:
[10.57938/a5211b15-b18b-4056-8580-f22ebe3d1685](https://doi.org/10.57938/a5211b15-b18b-4056-8580-f22ebe3d1685)

Published: 01/12/2018

Document Version:
Publisher's PDF, also known as Version of record

Document License:
Unspecified

[Link to publication](#)

Citation for published version (APA):
Hauzenberger, N., & Huber, F. (2018). *Model instability in predictive exchange rate regressions*. WU Vienna University of Economics and Business. Department of Economics Working Paper Series No. 276
<https://doi.org/10.57938/a5211b15-b18b-4056-8580-f22ebe3d1685>

Department of Economics
Working Paper No. 276

Model instability in predictive exchange rate regressions

Niko Hauzenberger
Florian Huber

December 2018



Model instability in predictive exchange rate regressions

Niko Hauzenberger^{*1,2} and Florian Huber²

¹WU Vienna University of Economics and Business

²Salzburg Centre of European Union Studies (SCEUS), Paris Lodron University of Salzburg

November 21, 2018

Abstract

In this paper we aim to improve existing empirical exchange rate models by accounting for uncertainty with respect to the underlying structural representation. Within a flexible Bayesian non-linear time series framework, our modeling approach assumes that different regimes are characterized by commonly used structural exchange rate models, with their evolution being driven by a Markov process. We assume a time-varying transition probability matrix with transition probabilities depending on a measure of the monetary policy stance of the central bank at the home and foreign country. We apply this model to a set of eight exchange rates against the US dollar. In a forecasting exercise, we show that model evidence varies over time and a model approach that takes this empirical evidence seriously yields improvements in accuracy of density forecasts for most currency pairs considered.

Keywords: Empirical exchange rate models, exchange rate fundamentals, Markov switching.

JEL Codes: C30, E32, E52, F31.

^{*}Corresponding author. Address: Vienna University of Economics and Business (WU), Welthandelsplatz 1, 1020 Wien, Austria. Email: niko.hauzenberger@wu.ac.at.

1 Introduction

Since the end of the Bretton Woods system in 1971, economists have been confronted with the challenging issue of designing empirical models of bilateral exchange rates which are also useful for forecasting applications. In a seminal contribution, [Meese and Rogoff \(1983\)](#) provide some early evidence that exchange rates are difficult to predict, at least in the short-run. Using a set of theoretical models in the spirit of [Frankel \(1979\)](#); [Dornbusch \(1976\)](#); [Hooper and Morton \(1982\)](#) that guide the choice of covariates included in the forecasting regression, [Meese and Rogoff \(1983\)](#) find that a simple random walk benchmark is difficult to outperform for most major exchange rate pairs. One reason for the dismal performance of most empirical and structural models is that, within a standard asset pricing framework, the high persistence of the underlying fundamentals in light of a discount factor near unity translates into highly persistent exchange rates. One consequence is that a random walk appears to be a benchmark extremely hard to beat (see [Engel and West, 2005](#)).

Over the years, a plethora of alternative econometric techniques emerged that provide more sophisticated means for analyzing exchange rate data to successfully improve longer-term predictions. The literature on unit roots and cointegration, for example, opened the way for tools to explicitly discriminate between short-term movements of a given currency pair and its long-run behavior. [Mark \(1995\)](#), for instance, applies an error correction model to a set of four exchange rates against the US dollar. Within this error correction framework, the exchange rate is assumed to return to its long-run equilibrium value determined by a simple monetary model eventually, with short-run fluctuations determined by own lagged values of the exchange rate and its fundamentals. The finding that exchange rates tend to be predictable in the medium- and long-run sparked a series of related contributions that corroborate this result for different periods and currency pairs ([Groen, 2000](#); [Mark and Sul, 2001](#); [Rapach and Wohar, 2002](#)).

More recently, several studies emphasized the usefulness of accounting for non-linearities in the underlying econometric models to provide more precise exchange rate predictions (see, for example, [Canova, 1993](#); [Sarno et al., 2004](#); [Mark, 2009](#); [Byrne et al., 2016](#); [Huber, 2016, 2017](#); [Huber and Zörner, 2018](#)). These non-linearities may relate to movements in the error variances of the models or to changes in the regression coefficients over time. [Byrne et al. \(2016\)](#) assess whether exchange rate predictions can be improved by using time-varying parameter models for a set of competing models with variable choice guided by a set of theoretical models.

The majority of the above literature deals with the question on whether a given empirical model that is loosely based on an underlying structural model outperforms

a set of competing models. However, another key source of non-linearities could stem from the fact that the underlying theoretical model changes over time, potentially jeopardizing the predictive fit of the econometric specification.¹ For instance, the recent success of Taylor rule based models (see [Engel and West, 2006](#); [Molodtsova et al., 2008](#); [Molodtsova and Papell, 2009](#); [Molodtsova et al., 2011](#)) can be attributed to the fact that the involved central banks did actually follow a policy rule that is closely related to a Taylor rule. With short-term interest rates, however, reaching the zero lower bound (ZLB) and central banks starting to adopt unconventional monetary policy measures, the question arises whether a Taylor rule still proves to be an adequate exchange rate model. In fact, recent literature on non-linear Taylor rules suggests that during the ZLB, Taylor rule based models lose their momentum against simple random walk specifications ([Byrne et al., 2016](#); [Huber, 2017](#)).

In this paper, we contribute to the literature by acknowledging this empirical evidence and propose a modeling framework that is capable of handling model instability over time in a flexible manner. We allow for dynamically switching between regimes, each incorporating different empirical exchange rate models. This is achieved by proposing a Markov switching (MS) regression model with each regime being characterized by different covariates arising from structural exchange rate models. In contrast to the existing literature, which relies on dynamic Bayesian model averaging techniques, our approach is an integrated modeling device. Through the introduction of time-varying transition probabilities, it allows to assess how the likelihood that a given model is adopted for each point in time depends on some signal variable. As signal variables, we adopt the (lagged) interest rate of the home and foreign country. This specification is motivated by the observation that Taylor rule fundamentals are good predictors in periods of the great moderation (with policy rates being significantly larger than zero), but are known for their weak performance after the Great Recession (characterized by policy rates close to zero).

We assess the merits of the proposed approach using a forecasting exercise for eight different exchange rates against the US dollar. By considering the resulting regime allocation and the transition probabilities, we examine whether structural models indeed tend to change and how this is related to movements in policy rates. The findings indicate that allowing for time-varying probabilities is a key feature, pointing towards a strong relationship between policy rates and the underlying transition distribution of the Markov process. In terms of forecasting, we find that our proposed model improves upon the random walk for selected currencies, both in terms of point and density predictions. The improvements for point forecasts are, however, muted. Comparing

¹For recent contributions that deal with this issue, see [Wright \(2008\)](#); [Beckmann and Schüssler \(2016\)](#); [Byrne et al. \(2018\)](#); [Beckmann et al. \(2018\)](#).

different model features reveals that a model based on a set of predictors that is based on fundamentals from various structural models combined with shrinkage priors and non-linearities (in the form of Markov switching) is also competitive.

The remainder of this paper is organized as follows. Section 2 discusses the four structural exchange rate models adopted while Section 3 proposes the econometric framework. The empirical application is presented in Section 4. Finally, the last section summarizes and concludes the paper. A technical appendix provides details on the estimation algorithm adopted.

2 Theoretical exchange rate models

In this section, we briefly discuss the main theoretical underpinnings to be used to guide covariate inclusion in the empirical model as well as to structurally identify the different regimes considered in our non-linear regression framework.

The point of departure for the discussion is a set of macroeconomic and financial quantities stored in an R -dimensional matrix \mathbf{X}_t ,

$$\mathbf{X}_t = \left(\underbrace{i_{t-1}, i_{t-1}^*, \pi_t, \pi_t^*, x_t, x_t^*, q_t}_{\mathbf{X}_{1t}}, \overbrace{m_t, m_t^*, y_t, y_t^*, e_t, p_t, p_t^*}_{\mathbf{X}_{2t}}, \underbrace{i_t, i_t^*}_{\mathbf{X}_{3t}} \right)', \quad (2.1)$$

with i_{t-1} denoting the lagged short-term interest rate, π_t inflation, x_t output gap, m_t money supply, y_t income, p_t price level, while the real exchange rate is denoted by q_t and the exchange rate by e_t .² The subsets of \mathbf{X}_t , \mathbf{X}_{jt} ($j = 1, \dots, 4$), represent the different structural models we are going to describe next.

A taxonomy of selected models of exchange rate determination

In the following, we provide a brief taxonomy of the theoretical models considered, depending on the theory adapted, guiding the specific partitions of \mathbf{X}_t .

- Our starting point is the model based on *Taylor rule fundamentals* (see [Molodtsova and Papell, 2009](#), for a recent forecasting study). This specification assumes that the set of predictors is given by \mathbf{X}_{1t} and thus includes the lagged short-term interest rate, inflation and the output gap of both the home and foreign country, and the real exchange rate. This model has proved to be successful in terms of describing exchange rate movements, both in-sample ([Engel and West, 2006](#)) and out-of-sample ([Molodtsova et al., 2008, 2011](#)). However, one critical assumption

²Asterisks denote foreign quantities. Moreover, y_t, m_t, p_t, q_t and e_t are measured in logarithms. For simplicity, we suppress subset-specific intercepts.

of this model is that the central bank at home and abroad is actively pursuing a Taylor rule-type monetary policy strategy. Especially during the recent period of the ZLB, this assumption could be violated, effectively leading to an inferior model fit.

- The second model considered is the *long-run monetary model*. The monetary model assumes that the covariates are given by \mathbf{X}_{2t} and include data on domestic and foreign money supply as well as the cross-country differences in income for a given income elasticity. As mentioned by [Rapach and Wohar \(2002\)](#), the long-run monetary model simply states that the price level of the home and foreign country is determined by the money supply and the level of production. Assuming purchasing power parities (PPP) and uncovered interest rate parity (UIP), one is able to relate the change of the exchange rate to supply and demand for money.
- Third, we consider a model based on PPP. This model is based on using \mathbf{X}_{3t} , leading to a regression model that includes domestic and foreign price indices. PPP originates from the theory of one price in goods markets, which in turn implies that the real exchange rate is supposed to revert to a long-run equilibrium level determined by relative prices. If this turns out to be true, the real exchange rate is a stationary process. However, [Sarno \(2005\)](#) highlights substantial persistence in real exchange rates. The convergence towards PPP is thus slow in the long run and real exchange rates typically display pronounced deviations from their PPP-implied fundamentals in the short run.
- Finally, we also augment our forecasting regression with the UIP model. By selecting \mathbf{X}_{4t} , this model simply establishes a relationship between the change in the exchange rate and the interest rate differential between home and abroad. Following [Chinn \(2006\)](#), UIP implies a positive one-to-one relationship between the interest rate differential and changes in the exchange rate. A positive change in the interest differential may be potentially followed by both an immediate and persistent appreciation in the short run, implying that UIP does not hold immediately. Here, we follow [Molodtsova and Papell \(2009\)](#), who address the UIP puzzle by not placing any restrictions on the coefficients.³

All these models have been shown to possess some merit in terms of predictive power. However, several recent studies find remarkable heterogeneity with respect to the fundamental model adopted (see, inter alia, [Wright, 2008](#); [Beckmann and Schüssler, 2016](#); [Byrne et al., 2018](#); [Beckmann et al., 2018](#)). In particular, [Beckmann et al. \(2018\)](#) argue

³See, for instance, [Engel \(2014\)](#); [Chinn \(2006\)](#) who observe coefficients that are less than one or even smaller than zero.

that an investor is capable to adjust his strategy by means of sequential learning and incorporating new information arriving in each point in time. The authors assume this behaviour is captured best by a dynamically changing set of fundamentals depending on historical forecast performance.

We now turn to describing our model framework that allows for dealing with issues of model instability in an intuitive way.

3 Controlling for model instability in empirical exchange rate models

In this section, we propose a model that controls for dynamic model instability by specifying a non-linear econometric framework. After summarizing the model structure in Subsection 3.1, we highlight the prior setup adopted in Subsection 3.2.

3.1 A Markov switching model specification

We now turn to describing the proposed Markov switching model with time-varying transition probabilities (MS-TVP). The key feature of our proposed framework is that, in general, it allows for switching between the fundamentals implied by K competing theoretical exchange rate models. We assume that exchange rate returns Δe_t follow an MS-TVP model given by

$$\Delta e_t = \mathbf{X}'_{S_t t-1} \boldsymbol{\beta}_{S_t} + \eta_t. \quad (3.1)$$

Hereby, $S_t \in \{1, \dots, K\}$ follows a first-order Markov process, $\boldsymbol{\beta}_k$ represents a vector of dimension M_k that collects the state-specific coefficients of state $S_t = k$ while $\eta_t \sim \mathcal{N}(0, \sigma_{S_t}^2)$ is a white noise shock with regime-specific variance $\sigma_{S_t}^2$. Note that, each $\boldsymbol{\beta}_k$ may exhibit different dimensions. We depart from the traditional literature on Markov switching models (see, among many others, [Hamilton, 1994](#); [Engel, 1994](#); [Filardo, 1994](#); [Amisano and Fagan, 2013](#); [Kaufmann, 2015](#); [Billio et al., 2016](#); [Huber and Fischer, 2018](#); [Casarin et al., 2018](#)) by assuming that the regimes are characterized by competing structural exchange rate models, implying that different fundamentals enter the predictive exchange rate regression at different points in time.

In the spirit of [Belmonte and Koop \(2014\)](#) and [Frühwirth-Schnatter \(2006\)](#) we introduce a selection matrix \mathbf{D}_{S_t} that entails switching between K alternative model specifications,

$$\Delta e_t = \mathbf{X}'_{t-1} \mathbf{D}_{S_t} \boldsymbol{\beta} + \eta_t, \quad (3.2)$$

with \mathbf{X}_{t-1} denoting an R -dimensional vector of the full set of economic fundamentals. We define $\boldsymbol{\beta} = (\boldsymbol{\beta}'_1, \dots, \boldsymbol{\beta}'_K)'$ as a stacked M -dimensional vector of regime-specific coefficients with $M = \sum_{j=1}^K M_j$, and $\boldsymbol{\sigma}^2 = (\sigma_1^2, \dots, \sigma_K^2)'$ collecting the K state-specific variances.

The selection matrix \mathbf{D}_k of state $S_t = k$ is a $R \times M$ dimensional matrix with binary indicators that allow for choosing β_k and \mathbf{X}_{kt-1} while zeroing out the elements in β and \mathbf{X}_{t-1} associated with the remaining models. For instance, we effectively obtain the model based on Taylor rule fundamentals, characterized through $S_t = 1$, by setting

$$\mathbf{D}_1 = \begin{pmatrix} \mathbf{I}_7 & \dots & \mathbf{0}_{7 \times 2} \\ \vdots & \ddots & \vdots \\ \mathbf{0}_{2 \times 7} & \dots & \mathbf{0}_{2 \times 2} \end{pmatrix}, \quad (3.3)$$

whereby $\mathbf{0}_{i \times j}$ is a $i \times j$ -dimensional matrix of zeros. Multiplying β from the left with \mathbf{D}_1 yields

$$\mathbf{D}_1 \beta = (\beta'_1, \mathbf{0}'_{5 \times 1}, \mathbf{0}'_{3 \times 1}, \mathbf{0}'_{2 \times 1})'. \quad (3.4)$$

From this discussion, it is clear that the matrix \mathbf{D}_{S_t} effectively controls the prevailing structural exchange rate model and the set of covariates to include in the state-specific regression. Notice, that an MS kitchen-sink regression is obtained by defining \mathbf{D}_{S_t} in such a way that in each states all economic indicators are included at all points in time.

Time-varying transition probabilities

Assuming constant transition probabilities is a standard (and potentially restrictive) assumption in Markov switching models (for economically motivated examples, see [Fildardo, 1994](#); [Amisano and Fagan, 2013](#); [Kaufmann, 2015](#)). Both, [Amisano and Fagan \(2013\)](#) and [Kaufmann \(2015\)](#) propose treating the transition distributions as being dependent on additional covariates. Here, and since our model features K regimes, we follow [Kaufmann \(2015\)](#) and parameterize the transition probabilities by a multinomial logit specification. Given the forecasting evidence provided in the literature quoted above, we assume that the transition probabilities depend on a measure of the monetary policy stance such as the policy rate. This captures the notion that if policy rates approach the ZLB, a Taylor rule-based model might become inadequate and the likelihood of a regime-shift could increase.

Let $\tilde{\mathbf{S}}_T = (S_1, \dots, S_T)'$ denote the full history of the state vector, the multinomial likelihood reads

$$P(S_t = j | S_{t-1} = k, \mathbf{Z}_t, \gamma) = p_{kj,t} = \frac{\exp(\mathbf{Z}'_t \gamma_{kj})}{1 + \sum_{l=1}^{K-1} \exp(\mathbf{Z}'_t \gamma_{kl})},$$

with category-specific regression coefficients $\gamma_{kj} = (\gamma_{0,kj}, \dots, \gamma_{N,kj})'$, collected in γ for all k and j . Moreover, we define \mathbf{Z}_t as an N -dimensional set of covariates. This set of

covariates is given by

$$\mathbf{Z}_t = (1, \mathbf{z}'_t, \mathcal{I}[S_{t-1} = 1], \dots, \mathcal{I}[S_{t-1} = K - 1])', \quad (3.5)$$

whereby \mathbf{z}_t is a vector of covariates that determine the dynamics of the transition probabilities while $\mathcal{I}(\bullet)$ denotes an indicator function that equals one if its argument is true. This implies that we capture a first-order Markov structure by including the previous states as additional regressors. Moreover, $\gamma_{0,kj}$ represents the intercept of the reference state $S_{t-1} = K$, and thus captures the corresponding time-invariant state persistence. Consistent with [Amisano and Fagan \(2013\)](#), we let the coefficients associated with \mathbf{z}_t be regime-invariant. It is worth noting that, if coefficients of \mathbf{z}_t are zero, we obtain a classic fixed transition probability Markov switching model. For further convenience, define $\tilde{\mathbf{Z}}_T = (\mathbf{Z}_1, \dots, \mathbf{Z}_T)$ and $\tilde{\mathbf{z}}_T = (\mathbf{z}_1, \dots, \mathbf{z}_T)$ to be the history of \mathbf{Z}_t and \mathbf{z}_t up to time T .

The specific choice of \mathbf{z}_t proves to be an important modeling decision. As mentioned above, our goal is to include a measure of the (conventional) monetary policy stance to signal a potential transition from Taylor rule-type based policy making to discretionary monetary policy actions such as quantitative easing (QE). In our case, we assume two early warning indicators $\mathbf{z}_t = (\tilde{i}_{t-1}, \tilde{i}_{t-1}^*)'$, the demeaned, lagged interest rate at home and abroad. Consequently, the three-dimensional coefficient vector γ_{kj} determines the sensitivity of the transition probability that drives the transition from the k th to the j th state. The demeaned covariates imply the centered parameterization of [Kaufmann \(2015\)](#), since a covariate $\tilde{\mathbf{z}}_T$ can be rewritten as a linear combination of its time-varying component $(\tilde{\mathbf{z}}_T - \bar{\mathbf{z}})$ and its mean $\bar{\mathbf{z}}$, which affects the time-invariant average state persistence. Demeaning covariates ensure that the time-invariant part does not depend on the scale of \mathbf{z}_t .

3.2 Prior specification and estimation strategy

Our approach is Bayesian and this implies that we have to carefully specify priors on the parameters of the model. Here, we follow [George and McCulloch \(1993, 1997\)](#) and specify a mixture of Gaussians prior on β_{ik} , the i th element of β_k . The prior is centered on the theoretically motivated restrictions, in order to test whether these restrictions hold true. The prior mean is stored in a M_k -dimensional vector $\underline{\beta}_k$ and summarized in [Table 1](#). A priori we assume a symmetric Taylor rule with same coefficients for the home and foreign country and do not consider interest smoothing (see [Molodtsova and Papell, 2009](#), for a detailed discussion). For the remaining models we center them on the implied long-run fundamental value.

Formally, this prior reads

$$\beta_{ik} | \delta_{ik} \sim \mathcal{N}(\underline{\beta}_{ik}, \tau_{ik,1}^2) \delta_{ik} + \mathcal{N}(\underline{\beta}_{ik}, \tau_{ik,0}^2) (1 - \delta_{ik}). \quad (3.6)$$

Here, we let $\tau_{ik,0}^2$ and $\tau_{ik,1}^2$ be prior variances (with $\tau_{ik,1}^2 \gg \tau_{ik,0}^2$), for $i = 1, \dots, M_k$, and $\underline{\beta}_{ik}$ denotes the i th element of $\underline{\beta}_k$. The first mixture component is referred to as the 'spike' component, tightly fixed around the prior mean $\underline{\beta}_{ik}$ while the second is called the 'slab' component, which translates into almost no prior influence. The indicator δ_{ik} serves to select the mixture component used. Following the semiautomatic approach of [George et al. \(2008\)](#), moreover, we scale the prior variances, $\tau_{ik,0}^2$ and $\tau_{ik,1}^2$, with variances of the ordinary least square estimates of the underlying structural model of state $S_t = k$.

This modeling approach constitutes a data-driven way of assessing whether coefficients should be pushed towards theoretically motivated restrictions or allowed to be closely related to the corresponding maximum likelihood estimate. Thus, if $\delta_{ik} = 0$, the posterior estimate of β_{ik} is strongly pushed towards the prior restriction $\underline{\beta}_{ik}$ while in the opposite case only little prior information on β_{ik} is introduced.

	Intercept	i_{t-1}	i_{t-1}^*	π_t	π_t^*	x_t	x_t^*	q_t	m_t	m_t^*	y_t	y_t^*	e_t	p_t	p_t^*	i_t	i_t^*
$\underline{\beta}_1$	0	0	0	1.5	-1.5	.5	-.5	0									
$\underline{\beta}_2$	0								1	-1	1	-1	-1				
$\underline{\beta}_3$	0												-1	1	-1		
$\underline{\beta}_4$	0															1	-1

Table 1: Prior mean $\underline{\beta}_k$ for each state

In what follows, we store all regime-specific indicators in a vector $\boldsymbol{\delta}_k = (\delta_{1k}, \dots, \delta_{M_k})'$ that corresponds to the block of β associated with the k th structural model. Each element of the latent variable $\boldsymbol{\delta}_k$ is a priori independently Bernoulli distributed,

$$\begin{aligned} p(\delta_{ik} = 1 | S_t = k) &= \underline{\omega}_{ik}, \\ p(\delta_{ik} = 0 | S_t = k) &= 1 - \underline{\omega}_{ik}, \end{aligned}$$

for hyperparameters $\underline{\omega}_{ik} \in \underline{\omega}_k$, an M_k -dimensional vector, chosen by the researcher. Again, as in the case of $\boldsymbol{\delta}_k$, the dimensionality of $\underline{\omega}_k$ and the elements directly corresponds to the coefficient vector β_k . A reasonable choice is $\underline{\omega}_{ik} = 0.5$, for all i, k , implying an equal prior probability of introducing significant prior information or using a relatively loose prior.⁴

⁴When considering a weakly informative coefficient prior, we define $\underline{\omega}_{ik}$ as being one for including all state-specific coefficients with certainty.

For the variances σ^2 , we assume an independent inverse Gamma prior for each element σ_i^2 ($i = 1, \dots, K$). More specifically, we set

$$\sigma_i^2 \sim \mathcal{G}^{-1}(a_0, A_0),$$

with a_0 and A_0 being scalars. The specific values for a_0 and A_0 are chosen to be weakly informative with hyperparameters $a_0 = 0.01$, $A_0 = 0.01$.

The prior distribution on the initial state is set to $p(S_0 = k) = 1/K$, for all k . (Kaufmann, 2015). Finally, for the coefficients of the multinomial logit model, we adopt a weakly informative and symmetric prior across all states. That is,

$$\gamma_{kj} \sim \mathcal{N}(\mathbf{0}, \underline{\mathbf{V}}),$$

for all k and $j = 1, \dots, K$ with $\underline{\mathbf{V}} = \zeta \mathbf{I}_K$, and ζ denoting a scalar. In the empirical application we set $\zeta = 100$.

In a Bayesian framework, we combine the likelihood with the prior to obtain the posterior distribution. In our case, the joint posterior density is intractable. Fortunately, however, the full conditional posterior distributions take simple forms, permitting Gibbs updating steps. The Markov chain Monte Carlo (MCMC) algorithm is described in more detail in Appendix A. In the empirical application, we repeat the algorithm 80,000 times, discard the first 30,000 draws as burn-in and define a thinning factor of ten, thus basing inference on 5,000 draws from the joint posterior.

Before proceeding to the empirical application, a brief word on identification is in order. Identification is necessary for structural interpretation of the states, but is not relevant if interest centers exclusively on the predictive density of the model (Frühwirth-Schnatter, 2001, 2006).⁵ Recall that in the present model, each regime is characterized by a different set of fundamentals, reflecting different theoretical exchange rate models. By exploiting the specific structure of the theoretical models we have imposed inequality constraints on the coefficients assuming which fundamentals enter each regime. The only potential source of non-identifiability occurs in case of more than one state pointing towards a random walk. However, pushing coefficients in direction of theoretical guided values is sufficient to disentangle regimes and fully identify the model. When considering the alternative specification, in which we always include all predictors, identification is certainly an issue. Moreover, each state is implicitly centered on a random walk a priori. In this case, we apply a permutation sampling step and solely focus on predictive densities.

⁵Markov switching models might suffer from identification problems due to the invariance of the likelihood with respect to permutations of the $K!$ possible labeling of the regimes, resulting in $K!$ modes.

4 Empirical application

This section starts by briefly describing the dataset and forecasting design adopted in Subsection 4.1. Then, we discuss key in-sample features of the model in Subsection 4.2. Finally, Subsection 4.3 presents the main forecasting results, discriminating between point and density forecasting performance of all models considered.

4.1 Data, forecasting design, and competing models

In this paper, our aim is to forecast bilateral exchange rates for Australia, Canada, Japan, Norway, South Korea, Sweden, Switzerland and the United Kingdom relative to the US dollar. We collect monthly data on nominal exchange rates, industrial productions, monetary aggregates, three-month money market rates and consumer price indices for countries under consideration. Table 2 depicts the data transformations of variables. Table 3 provides an overview of the data alongside information on time coverage and source of economic fundamentals.

Variable	Description	Transformation	Comments
EXR	Nominal exchange rate	log difference	—
IP	Industrial production	log	—
M	Money aggregate	log	—
3M-IR	3M Money market rate	—	—
CPI	Consumer price index	log	—
INF	Inflation	log differences of CPI	—
REXR	Real exchange rate	$\log(\text{EXR}) + \log(\text{CPI}^*) - \log(\text{CPI})$	—
IP-GAP	Output gap	HP filter	$\lambda = 14400$ for monthly data

Table 2: Transformation of variables.

Country	Coverage	EXR	IP (2010 = 100)	M	3M-IR,	CPI (2010 = 100)
Australia (AU)	1975M06:2017M09	IFS	OECD	M1, OECD	OECD	OECD
Canada (CA)	1973M01:2017M09	IFS	OECD	M1, OECD	OECD	OECD
Japan (JP)	1973M01:2017M03	IFS	OECD	M1, OECD	IFS	OECD
Norway (NO)	1979M01:2017M09	IFS	OECD	M1, OECD	OECD	OECD
South Korea (KR)	1981M01:2017M09	IFS	IFS	M1, OECD	IFS	OECD
Sweden (SE)	1973M01:2017M06	IFS	OECD	M3, OECD	IFS	OECD
Switzerland (CH)	1974M01:2017M09	IFS	OECD	M1, OECD	OECD	OECD
United Kingdom (UK)	1973M01:2017M02	IFS	OECD	M0, IMF/FRED	IFS	OECD
United States (US)	1973M01:2017M09		OECD	M1, OECD	OECD	OECD

Note: All quantities are seasonally adjusted, except EXR and 3M-IR. IP of Australia and Switzerland are interpolated to monthly frequency with cubic spline.

Table 3: Sources of economic fundamentals.

In order to assess whether time-varying transition probabilities improve predictive accuracy, the proposed model framework is benchmarked with MS specifications with fixed transition probabilities (labeled MS-FT), as well as standard structural exchange

rate models that are estimated under weakly informative priors (labeled linear). These linear benchmarks are based on Taylor rule, monetary, PPP, and UIP fundamentals. Therefore, in the forecasting exercise, the set of competing models is divided in the three overall classes: MS-TVP, MS-FT and linear. Moreover, we consider not only theoretically motivated MS-TVP and MS-FT specifications, but also models that include all macroeconomic indicators of \mathbf{X}_{t-1} within each state (labeled kitchen-sink). For the kitchen-sink regressions, we consider different numbers of states, ranging from two to four regimes. Moreover, to allow for state-specific shrinkage in kitchen-sink regressions, the SSVS prior described above is centered on zero and different state-specific indicators are estimated. To assess the role of allowing for heteroscedasticity in forecasting exchange rates, we also consider MS-TVP specifications with state-specific variances. All models are then benchmarked to the random walk without drift.

We evaluate predictive accuracy by means of a recursive pseudo out-of-sample forecasting exercise. This implies choosing a initial estimation period that ranges from $t = 1$ up to $t = T_0$ with the remaining periods used as a hold-out sample. In the present application, we estimate all models using data up to 2004M12 and then proceed by computing h -step-ahead predictions for $t = T_0 + 1$. After obtaining draws from the corresponding predictive distributions, we consequently expand the initial estimation period by one month. This procedure is repeated until the end of the sample is reached.

To rank forecasts, we rely on cumulative squared forecast errors (CSFEs) to assess the quality of point forecasts. As point predictions, we take the posterior median of the predictive density. Turning to density forecasts, we follow [Geweke and Amisano \(2010\)](#) and rely on the log predictive score (LPS) to measure density forecasting accuracy. This has the advantage that, conditional on the proposed model and data, uncertainty surrounding the parameters and latent quantities is integrated out. After obtaining the LPS, we compute log predictive Bayes factors (LBFs) for the entire hold-out sample by computing the difference between the LPS of a given model relative to the random walk.

4.2 Inspecting evidence for model instability

We now turn to assess whether our model proposed in Subsection 3.1 signals significant shifts in the underlying structural representation. Figure 1 summarizes the mean of the filtered state probabilities for the eight exchange rates considered. In general, we observe that the regime dynamics across countries share one common feature. The models based on Taylor rule (state 0) and the UIP fundamentals (state 4) appear to be the dominant states before the global financial crisis in 2008/2009. After that period, however, model evidence changes significantly for the majority of countries. More

precisely, models based on monetary (state 2) and PPP (state 3) fundamentals tend to receive more posterior support.

Compared to the remaining currencies, the Swiss franc (see Fig. 1(c)) exhibits a somewhat higher regime-switching frequency. Figure 1, moreover, suggests that hitting the ZLB does shift filtered probabilities, supporting regimes other than Taylor rule fundamentals (state 0) for countries such as Australia, Canada, South Korea, Sweden, and the United Kingdom.⁶ Countries such as Japan and Switzerland, on the other hand, indicate an opposite dynamic, namely a shift of probabilities towards the Taylor rule state.

Taking a closer look at the United Kingdom, Taylor rule fundamentals are the predominant regime, reflecting the fact that these quantities tend to describe exchange rates well in times when the primary policy rule of the Bank of England is the Taylor rule. After this period, transition probabilities point towards a first shift during the crisis of the European Monetary System. After the financial crisis, and upon hitting the ZLB, the model based on Taylor rule fundamentals receives only limited posterior support. It is noteworthy that after 2010, the short-term interest rate (both at home and abroad) is stuck at zero (and almost constant). This implies that the model based on interest fundamentals closely mimic a random walk with drift during this period, even without introducing shrinkage.

The transition probabilities, depicted in Figs. 2 and 3, generally track the movements in filtered state probabilities, providing considerable evidence of time-varying transition distributions. Our findings thus suggest that a measure of the monetary policy stance at home and abroad tends to drive transitions between structural models. This is consistent with our conjecture that during the period of the ZLB, using Taylor rule-based exchange rate models might be inappropriate, at least from an in-sample perspective.

4.3 Forecasting results

In this section, interest centers on the predictive performance of our proposed MS-TVP specification. The discussion in the last section highlights that it proves to be important to allow for time-varying transition probabilities in-sample for several exchange rate pairs. This suggests that parameterizing the transition distributions with additional covariates helps to avoid situations where the model gets stuck within a certain state. Amisano and Fagan (2013) and Kaufmann (2015) highlight this issue and point towards advantages of explaining the regime-switching behavior of the model as opposed to a model based on constant transition probabilities. The key question, however,

⁶Notice that Australia never hit the ZLB during the sample, but the foreign country the United States did. Moreover, it is worth emphasizing that Japan already hit the ZLB in the midst of the 1990s.

is whether this additional flexibility also improves predictive performance. We answer this question using both, point and density forecasts.

Point forecasts

Figures 4 and 5 present the evolution of CSFEs of one-step-ahead forecasts for the best performing models across the considered model classes. We consider all linear predictive exchange rate regressions and the five best performing MS-TVP and MS-FT models, according to the CSFEs at the end of the hold-out-sample. CSFEs of all models are presented relative to the CSFEs of the random walk benchmark. Thus, values below zero indicate more accurate forecast relative to random walk predictions. Here, we focus on one-step-ahead forecasts since we find that models that perform well at the one-step-ahead horizon also do well for $h > 1$ periods ahead.⁷ When considering density forecasts, we report the results for higher order predictions as well.

Turning to the actual results, we observe pronounced differences across countries. For instance, in Australia, Canada, Norway, Sweden and the United Kingdom, modeling non-linearities pays off, in particular during periods of financial turmoil, outperforming forecasts of linear models as well as the random walk benchmark. herefore, one interesting finding is that controlling for heteroscedasticity also tends to exert a positive effect on the point forecasting performance during volatile periods of the business cycle.

By contrast, for South Korea and Switzerland, the random walk appears to be hard to beat. In general, we observe error ratios that are close to unity when averaged over the full hold-out sample. This indicates that including more information does not necessarily translate into improved point predictions relative to a simple no-change forecast for these two economies. Again, we find some heterogeneity in relative forecasting performance over time.

Turning to the performance of the theoretically inspired MS-TVP specifications, we observe strong forecasting accuracy for Japan, Norway, South Korea, and the United Kingdom, at least for one specification (marked red in Figs. 6 and 7). On the other hand, it appears that kitchen-sink specifications dominate theoretically inspired regimes for Australia, Canada, Switzerland and Sweden. In particular, this holds true for Switzerland as indicated by the absence of a red colored line in Fig. 6(c) for MS-TVPs.

When comparing MS-TVP to MS-FT models, we observe that MS-FT specifications appear to be more robust over time in terms of CSFEs. This can be seen by noting that the forecast errors of the MS-FTs feature fewer outliers. In general, we observe a better performance of MS-TVP models for Australia, Japan and Norway, at least for

⁷Additional results for $h > 1$ step ahead forecasts are provided on request.

one model, compared to MS-FT models; although MS-FT models perform well for the United Kingdom and Canada.

Moreover, accuracy differences across simple linear specifications appears to be diverse. For the United Kingdom, Sweden and Japan, we observe an inferior predictive performance for at least two linear models. In particular, models based on monetary fundamentals exhibit a weak forecast performance relative to the remaining models under scrutiny. For Australia, Canada and South Korea, all linear models do well and show a similar point forecast performance as the random walk. When focusing on Taylor rule fundamentals, the linear regression performs well at the beginning of the hold-out-sample (see, for example, in the case of the United Kingdom), but exhibits a systematic loss of performance in periods after the financial crisis. This can be explained by the fact that Taylor rule based models build on the assumption that both central banks' monetary policy might be well described by a Taylor rule. However, after the financial crisis, interest rates hit the ZLB and central banks increasingly adopted non-standard policy measures. This, in turn, leads to a deteriorating performance of this model class, effectively confirming findings reported in the recent literature (see, for example, [Byrne et al., 2016](#); [Molodtsova and Papell, 2012](#)).

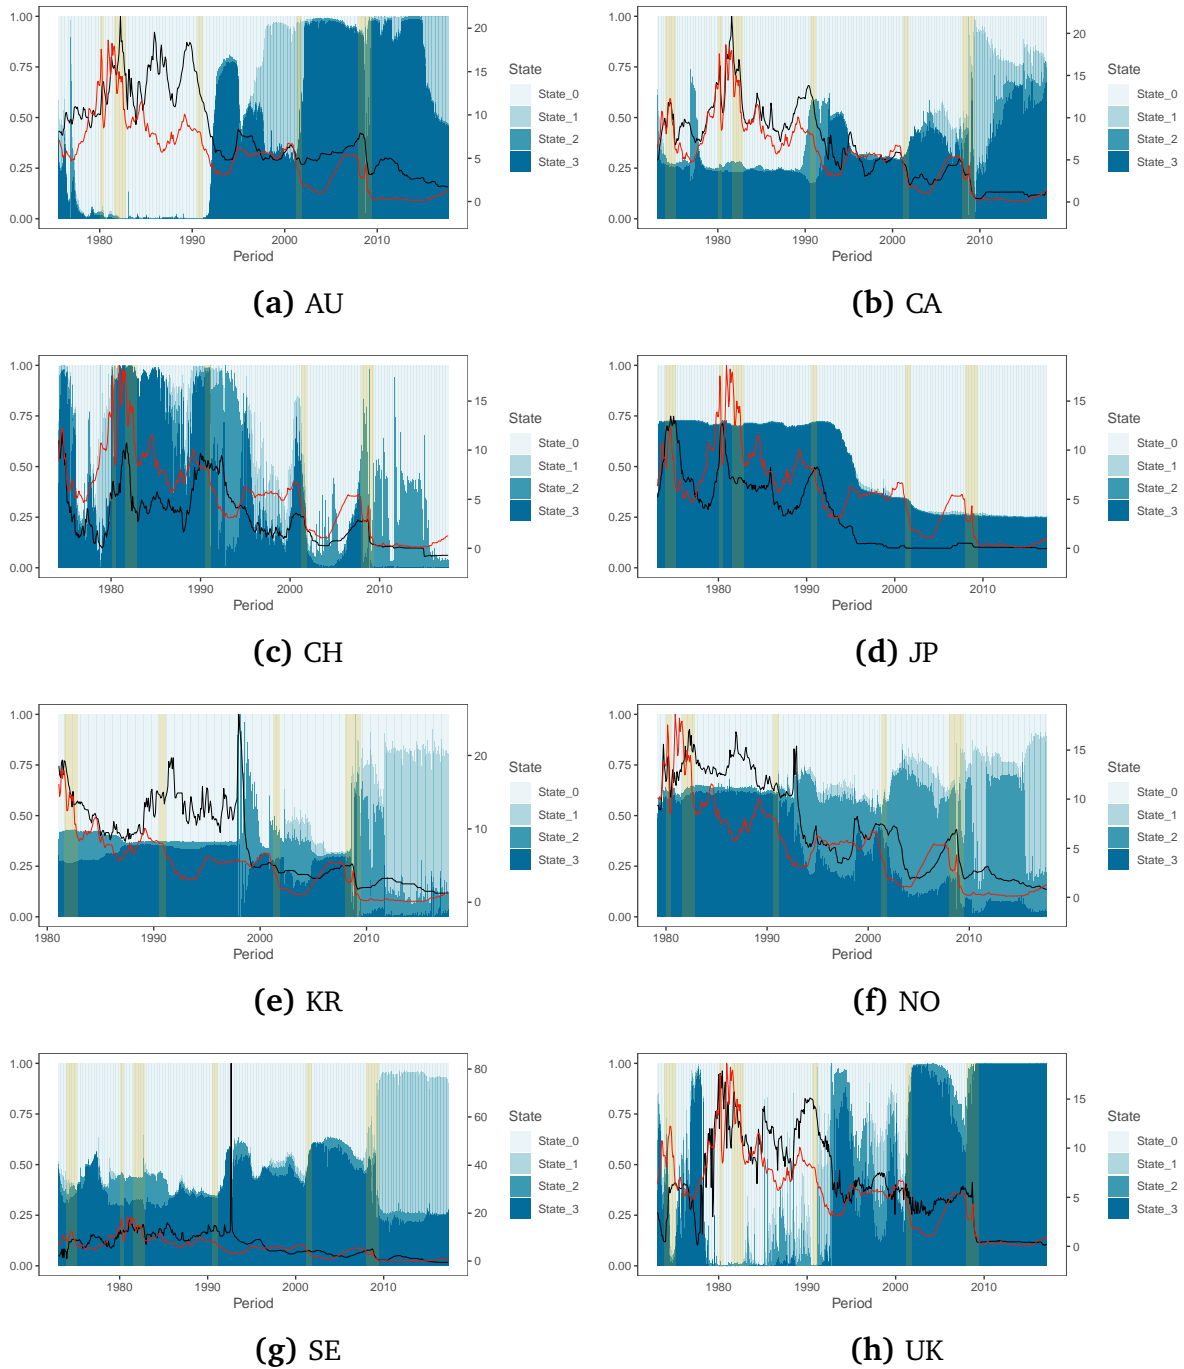


Fig. 1: Mean posterior state probabilities for a weakly informative prior and a common variance across states. State 0 indicates Taylor rule fundamentals, state 1 monetary fundamentals, state 2 PPP fundamentals and state 4 UIP fundamentals. The vertical bars (yellow) indicate NBER recessions for the US. The black solid line depicts the country-specific interest rate and the red solid line the interest rate for the US. The left-axis shows the probabilities and the right-scale the values of interest rates.

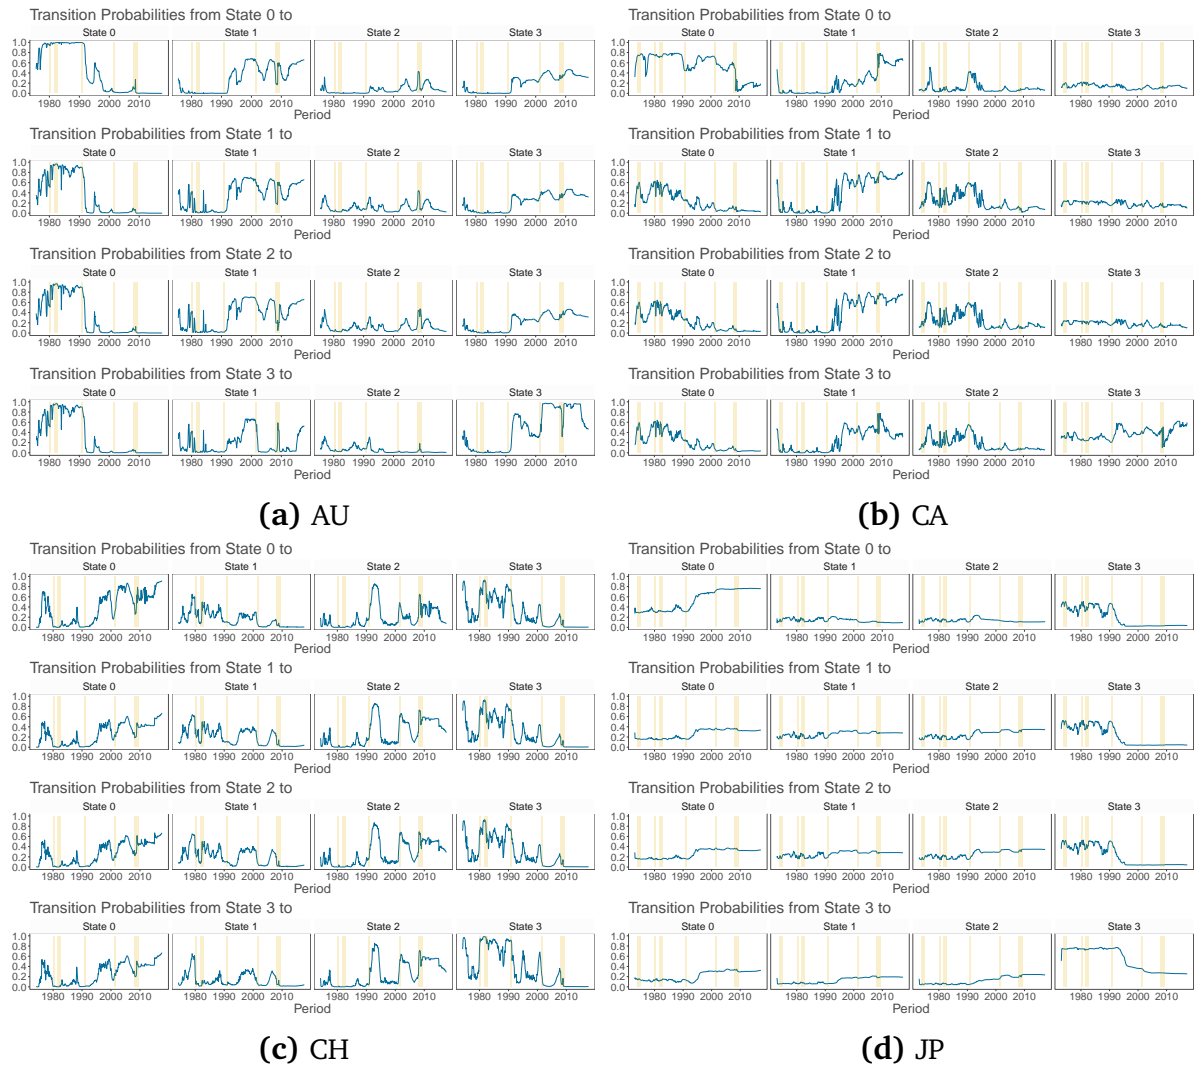


Fig. 2: The blue line depicts the posterior mean of time-varying transition probabilities for each state with a weakly informative prior and common variance across states. State 0 indicates Taylor rule fundamentals, state 1 monetary fundamentals, state 2 PPP fundamentals and state 4 interest rate fundamentals. The vertical bars (yellow) indicate NBER recessions for the US.

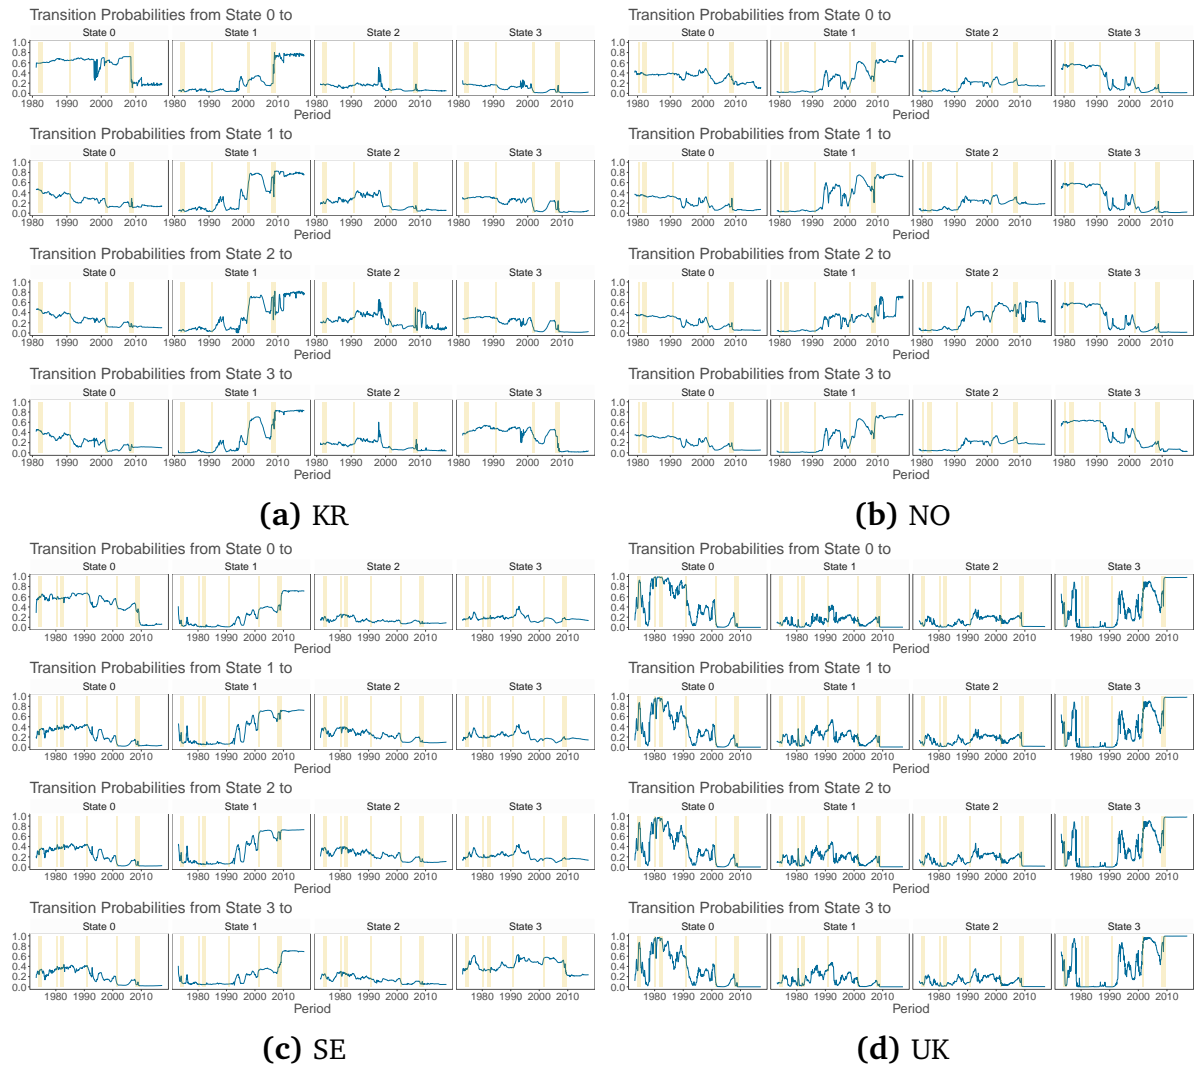
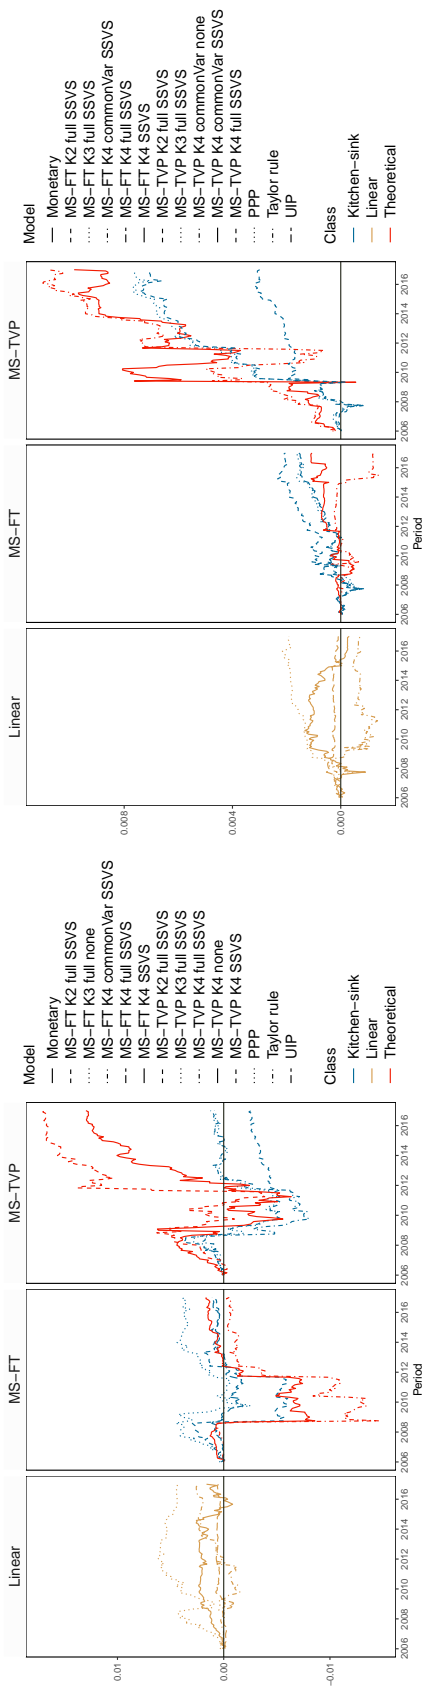
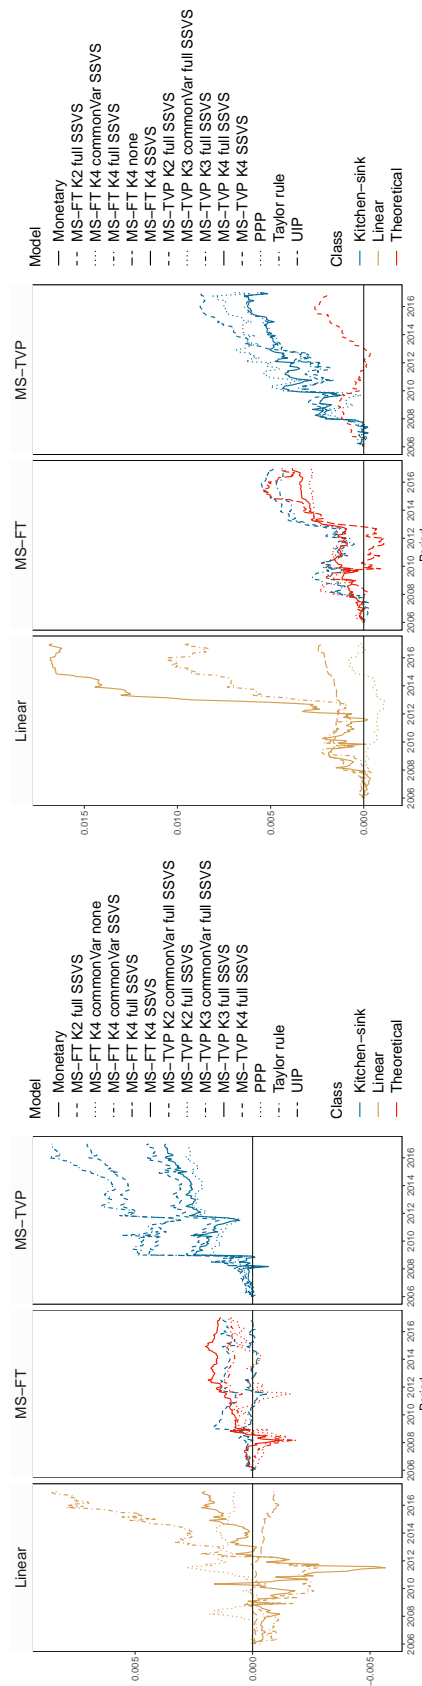


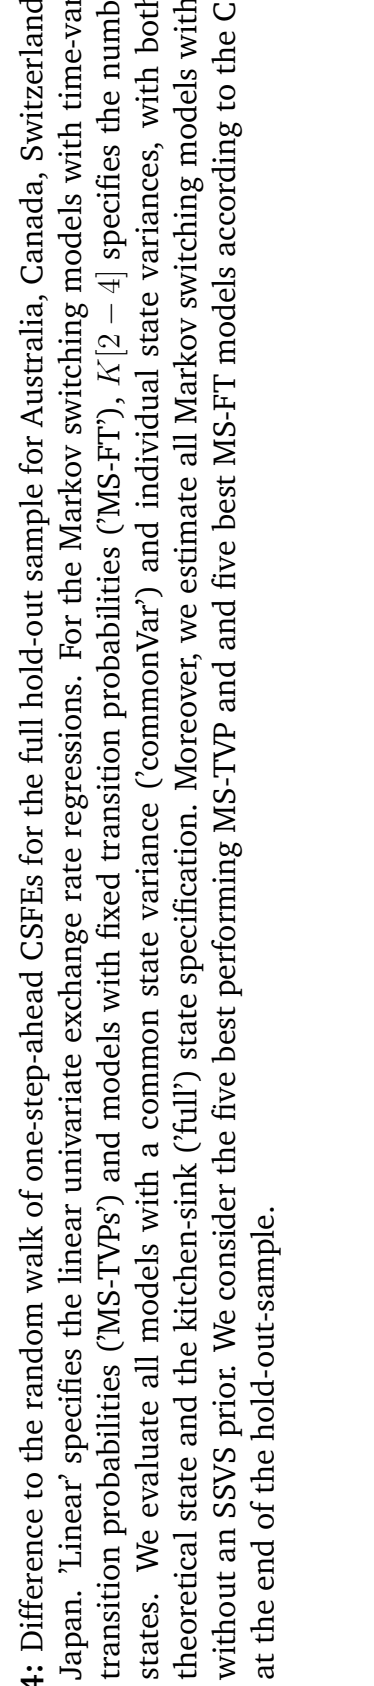
Fig. 3: The blue line depicts the posterior mean of time-varying transition probabilities for each state with a weakly informative prior and common variance across states. State 0 indicates Taylor rule fundamentals, state 1 monetary fundamentals, state 2 PPP fundamentals and state 4 interest rate fundamentals. The vertical bars (yellow) indicate NBER recessions for the US.



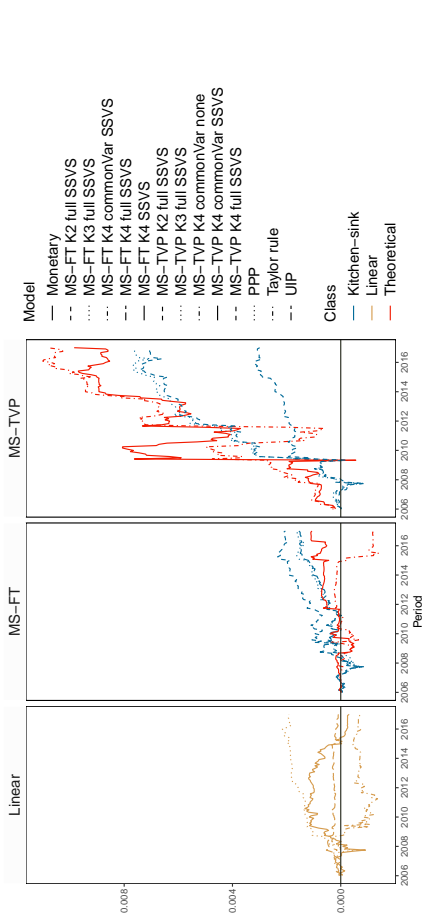
(a) AU



(b) CA



(c) CH



(d) JP

Fig. 4: Difference to the random walk of one-step-ahead CSFEs for the full hold-out sample for Australia, Canada, Switzerland and Japan. 'Linear' specifies the linear univariate exchange rate regressions. For the Markov switching models with time-varying transition probabilities ('MS-FT'), $K[2 - 4]$ specifies the number of states. We evaluate all models with a common state variance ('commonVar') and individual state variances, with both the theoretical state and the kitchen-sink ('full') state specification. Moreover, we estimate all Markov switching models with and without an SSVS prior. We consider the five best performing MS-TVP and five best MS-FT models according to the CSFEs at the end of the hold-out-sample.

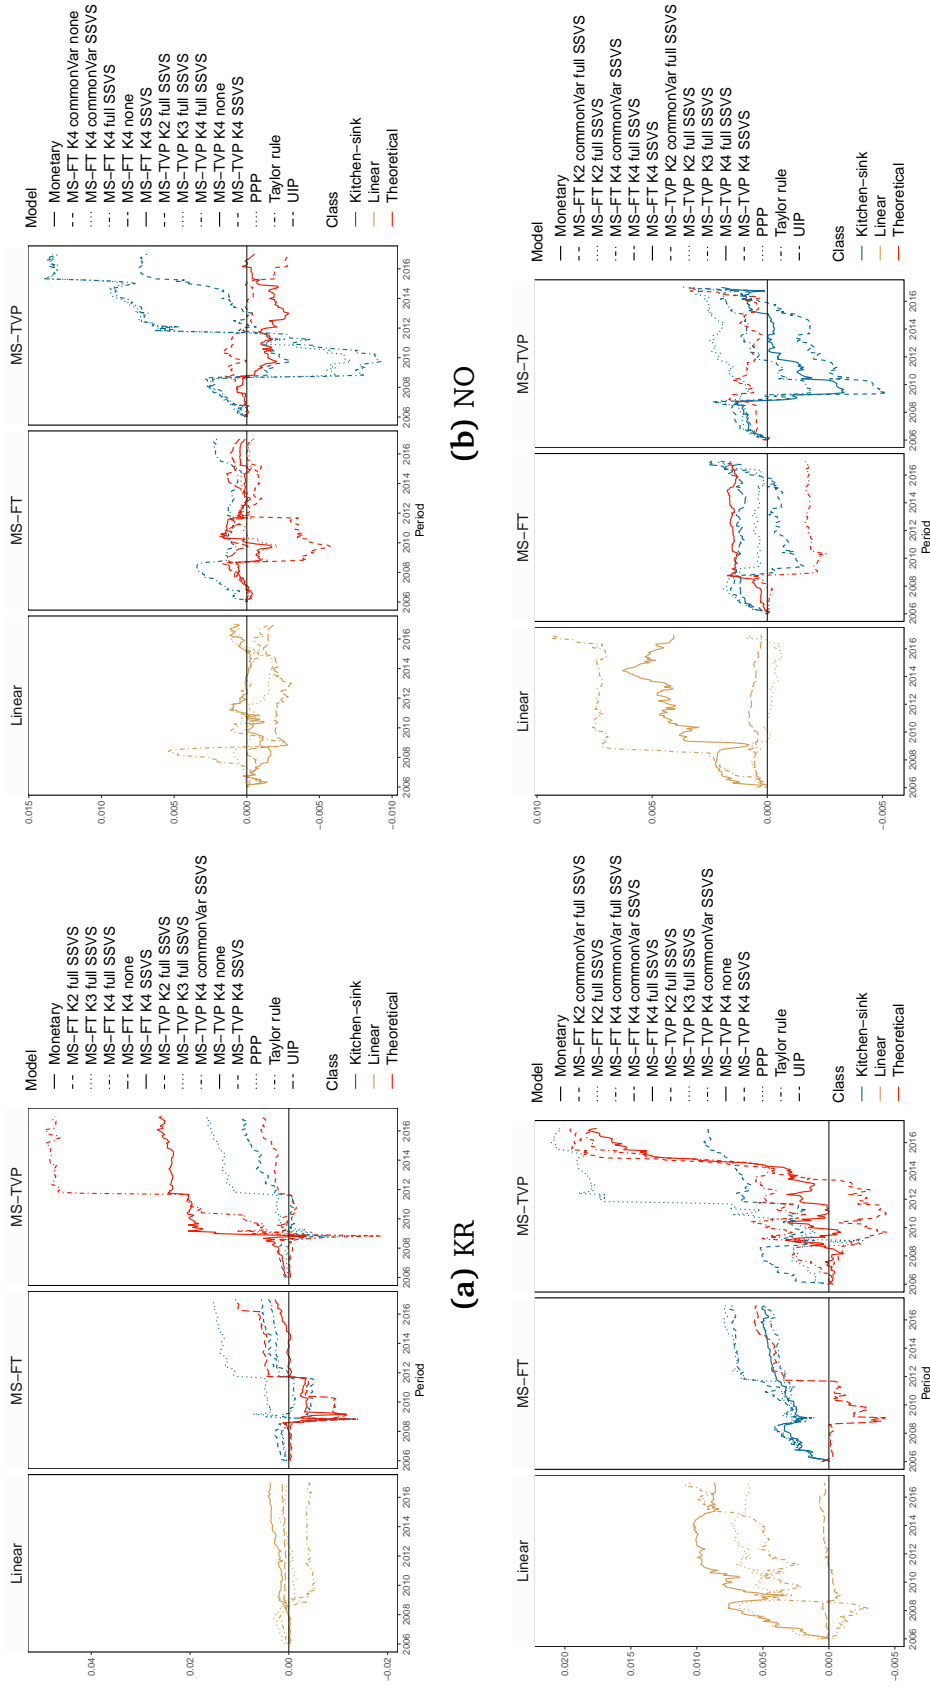


Fig. 5: Difference to the random walk of one-step-ahead CSFEs for the full hold-out sample for South Korea, Norway, Sweden and the United Kingdom. 'Linear' specifies the linear univariate exchange rate regressions. For the Markov switching model with time-varying transition probabilities ('MS-TVPs') and models with fixed transition probabilities ('MS-FT'), $K[2-4]$ specifies the number of states. We evaluate all models with a common state variance ('commonVar') and individual state variances, with both the theoretical state and the kitchen-sink ('full') state specification. Moreover, we estimate all Markov switching models with and without an SSVS prior. We consider the five best performing MS-TVP and five best MS-FT models according to the CSFEs at the end of the hold-out-sample.

Density forecasts

Tables 4 and 5 depict a summary of all models' LBFs for all currency pairs considered. Values highlighted in green point towards outperformance of a model relative to the random walk while red values signal a weaker predictive performance when benchmarked against the random walk. To provide a dynamic picture of LBFs over time, Figs 6 and 7, again, show the LBFs of all linear predictive exchange rate regressions and the five best performing MS-TVP and MS-FT models.

In general, both tables attest non-linear specifications good predictive power, while linear models display a somewhat weaker forecast performance. Furthermore, our results suggest that non-linear models that perform well in terms of point predictions also exhibit high predictive capabilities in terms of density forecasts. However, we find that predictive performance evolves differently for CSFEs and the LBFs. Density forecasts strengthen the argument in favor of non-linear models, as the performance gains of the MS models are even more sizable in periods of high exchange rate volatility while accuracy losses in tranquil times are rather muted.

Although we do not observe a single dominant non-linear model across forecast horizons and countries, Tables 4 and 5 suggest that at least one non-linear specification outperforms the random walk and the linear counterparts as well. One exception proves to be Japan, for which the best specification is either the random walk (for one-step ahead predictions) or the linear PPP model (for longer horizons).

For all forecast horizons considered, non-linear specifications do well for Australia, Canada, South Korea, Sweden and the United Kingdom. Specifically, the MS-TVP kitchen-sink specification with four states coupled with state-specific variances, constituting the most flexible specification, has good predictive power for Australia, Canada and South Korea. For Australia, this model is the single best performing model across forecast horizons and for South Korea, it is the best for three- and twelve-step-ahead forecast and the second best specification in terms of one-step ahead forecasts. Moreover, as shown in Figs. 6 and 7, the poor performance of linear exchange rate models is even more pronounced for density forecasts. In particular, the linear structural regression display a sharp decline in predictive power after the financial crisis, corroborating findings in [Byrne et al. \(2016\)](#); [Molodtsova and Papell \(2012\)](#).

Turning to the question whether allowing for heteroscedasticity pays off in terms of density forecasting, we find substantial evidence that this additional flexibility proves to be important. The gains in predictive accuracy of non-linear models can mainly be attributed to the more flexible variance specification of MS models. This can also be seen by comparing MS specifications with a common variance across states with MS models that feature individual state variances. For these models, we observe a slight

accuracy premium relative to their homoscedastic counterparts. Allowing for state-specific variances thus appears to be an important ingredient of a successful forecasting model. However, this increased flexibility comes at a cost. Specifically, we observe that during normal periods, relative predictive accuracy declines steadily for several MS models (see, for example, Fig. 7(c)), in line with recent evidence provided in [Abbate and Marcellino \(2018\)](#).

Contrasting MS-TVP with MS-FT models, Figs. 6 and 7 show that MS-TVP models yield more precise density forecasts for Australia, Norway and South Korea and Switzerland, while yielding an almost equivalent performance for Canada. Moreover, with time-varying transition probabilities, theoretically motivated specifications play a more important role than for MS-FT models. This points towards potential accuracy premia obtained by allowing for time-varying transition probabilities and thus sharpen inference surrounding the regime allocation.

Theoretically motivated MS-TVP specifications exhibit good forecast performance for Australia, Canada and Norway. Considering the results in Sweden, theoretically motivated MS-TVPs perform well during prolonged periods of high exchange rate volatility. In periods of low volatility, however, this specification is slightly outperformed by competing specifications. By contrast, we find that for South Korea, MS-TVP kitchen-sink regressions improve upon our proposed MS-TVP model that allow for switching across structural exchange rate regressions. For the remaining countries in Figs. 6 and 7, no clear pattern emerges when comparing both model types. For example, using structural MS-TVP specifications yields strong increases in predictive power during the financial crisis but a weaker performance afterwards. For kitchen-sink MS-TVP regressions, we find no gain during the crisis but, at the same time, no subsequent loss in the aftermath of the crisis.

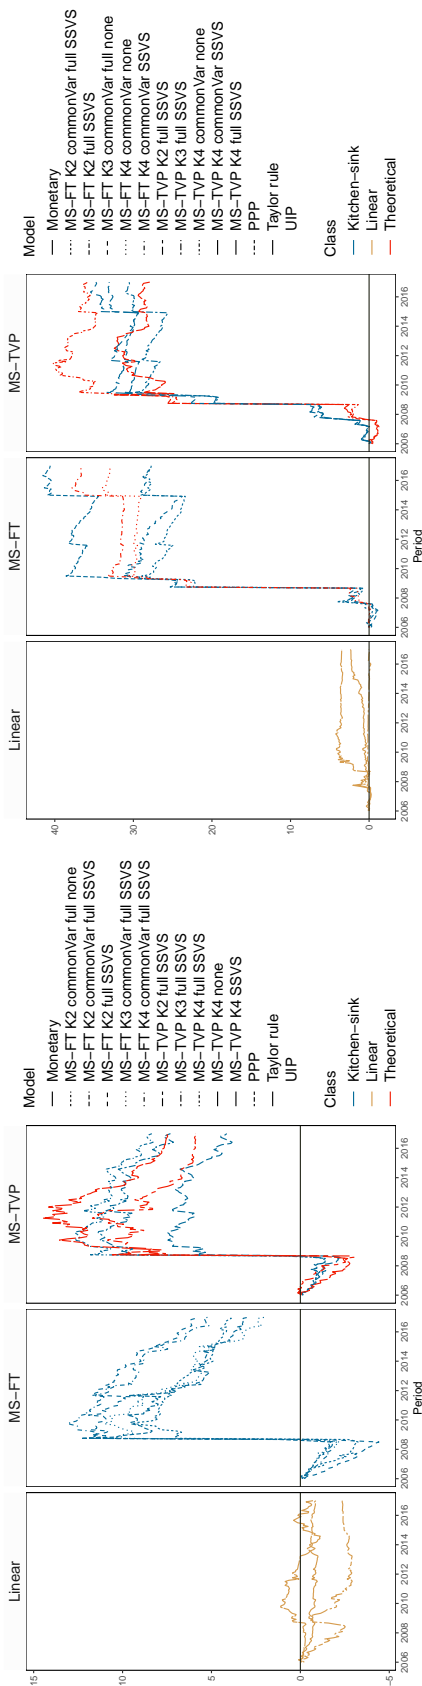
Finally, we assess whether using shrinkage priors on the coefficients improves forecasts. Tables 4 and 5 indicate that shrinkage generally translates into better results in pairwise comparisons with the corresponding non-shrinkage counterpart. This observation, however, is not consistent across all models, countries and forecast horizons considered. In particular, using a kitchen-sink regression without shrinkage leads to poor forecast performance, as already shown by [Li et al. \(2015\)](#); [Wright \(2008\)](#). Turning to theoretically motivated MS models provides mixed insights on whether using shrinkage is useful. We conjecture that this stems from the fact that adopting Markov switching specifications with theoretically defined regimes already introduces a certain amount of regularization that helps avoid overfitting.

	Specification			Log predictive Bayes factor						
	States	Shrinkage	Variance	1-step		3-step		12-step		
				AU			CA			
MS-TVP THEORETICAL	K = 4	None	Common	-16.208	-11.860	-7.476	35.802	28.588	22.695	
		None	State-specific	7.449	-0.068	-0.953	1.911	6.161	28.652	
		SSVS	Common	-6.908	-3.737	2.113	27.912	24.035	27.860	
		SSVS	State-specific	5.904	-0.597	4.272	3.577	23.006	10.700	
MS-FT THEORETICAL	K = 4	None	Common	-1.941	-2.809	-1.232	32.992	29.238	29.081	
		None	State-specific	1.693	1.070	-2.181	8.760	7.406	1.590	
		SSVS	Common	1.696	0.602	0.386	36.737	28.763	30.674	
		SSVS	State-specific	1.261	0.144	-2.762	7.856	9.720	7.995	
MS-TVP KITCHEN-SINK	K = 2	SSVS	Common	-22.236	-12.630	-26.266	-42.757	-38.942	-43.791	
		SSVS	State-specific	4.142	5.954	2.386	33.067	33.136	31.270	
	K = 3	SSVS	Common	-19.408	-19.043	-24.598	-47.925	-47.945	-50.501	
		SSVS	State-specific	7.306	6.406	5.417	30.543	31.877	24.891	
	K = 4	SSVS	Common	-23.884	-29.087	-29.255	-50.257	-57.303	-53.021	
		SSVS	State-specific	8.342	7.407	6.263	34.765	35.336	29.015	
	MS-FT KITCHEN-SINK	K = 2	SSVS	Common	2.911	0.250	0.201	27.733	27.812	24.748
			SSVS	State-specific	5.765	4.801	3.528	40.572	39.997	37.250
K = 3		SSVS	Common	1.962	3.487	1.319	-31.072	-25.389	-23.545	
		SSVS	State-specific	1.856	1.061	0.368	25.606	24.982	22.641	
K = 4		SSVS	Common	5.192	3.376	3.624	24.224	27.691	19.228	
		SSVS	State-specific	-16.540	-17.012	-17.715	15.123	14.930	12.654	
LINEAR	Taylor rule			-0.866	-1.661	-0.592	3.496	3.324	-0.164	
	Monetary			-0.709	-0.052	-1.914	2.379	1.594	0.007	
	PPP			-2.362	-2.061	-3.797	0.050	0.347	-1.133	
	UIP			-0.389	-0.286	-0.580	-0.479	0.478	0.207	
				CH			JP			
MS-TVP THEORETICAL	K = 4	None	Common	-6.457	-4.135	-2.080	-11.153	-10.498	-6.965	
		None	State-specific	0.081	-3.537	0.655	-5.688	-7.339	-1.672	
		SSVS	Common	-4.213	-6.261	-0.866	-5.443	-4.734	-0.660	
		SSVS	State-specific	0.176	-2.195	-1.363	-1.258	-3.183	-0.842	
MS-FT THEORETICAL	K = 4	None	Common	0.210	-1.905	-2.192	-2.310	-3.072	-0.849	
		None	State-specific	-1.142	-1.636	-1.787	-1.599	-1.536	-3.688	
		SSVS	Common	-0.136	-0.471	-2.668	-0.151	-1.152	-1.759	
		SSVS	State-specific	-0.385	-0.431	-3.329	-1.722	-0.936	-1.896	
MS-TVP KITCHEN-SINK	K = 2	SSVS	Common	0.497	-2.012	-2.175	-9.272	-2.967	-1.222	
		SSVS	State-specific	1.353	0.408	-0.569	-2.964	-2.992	-0.679	
	K = 3	SSVS	Common	-0.213	-2.987	-4.000	-1.906	-2.174	0.331	
		SSVS	State-specific	0.619	-0.520	-1.308	-2.411	-2.367	-1.370	
	K = 4	SSVS	Common	-1.158	-5.658	-5.427	-3.400	-3.803	-1.843	
		SSVS	State-specific	-0.543	-1.922	-2.322	-3.669	-3.662	-3.095	
	MS-FT KITCHEN-SINK	K = 2	SSVS	Common	-2.365	-2.906	-2.920	-2.230	-3.487	-2.038
			SSVS	State-specific	-5.686	-6.616	-6.344	-6.240	-8.560	-8.543
K = 3		SSVS	Common	-0.256	-3.726	-4.255	-3.233	-4.880	-3.300	
		SSVS	State-specific	-6.504	-9.566	-10.908	-12.533	-13.159	-13.005	
K = 4		SSVS	Common	-1.772	-3.870	-4.908	-8.006	-7.445	-6.180	
		SSVS	State-specific	-27.363	-27.798	-27.909	-27.305	-26.146	-25.993	
LINEAR	Taylor rule			-3.356	-2.715	-4.107	-4.302	-3.359	-3.218	
	Monetary			-0.617	-2.009	-3.243	-7.684	-6.439	-1.597	
	PPP			-0.787	-1.344	-1.349	-0.369	0.193	0.769	
	UIP			0.090	0.002	0.035	-0.507	-0.548	-0.559	

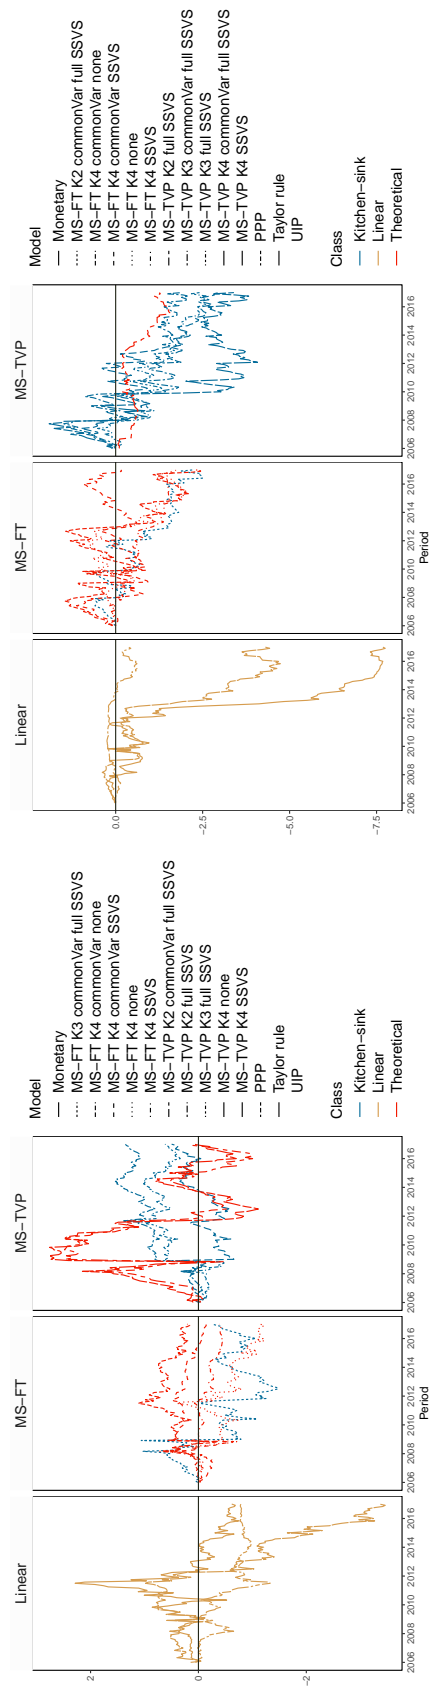
Table 4: Cumulative one-, three-, and twelve-step-ahead LBFs (random walk benchmark) at the end of the full hold-out sample summarized for Australia, Canada, Switzerland and Japan. Values highlighted green are greater than zero, values highlighted red are smaller than zero, indicating a better or a weaker performance compared to the random walk. Best model in bold.

	Specification			Log predictive Bayes factor						
	States	Shrinkage	Variance	1-step		3-step		12-step		
				KR			NO			
MS-TVP THEORETICAL	K = 4	None	Common	-91.338	-73.679	-40.795	-3.684	-2.929	-0.518	
		None	State-specific	15.820	8.041	5.458	5.623	-0.229	-0.430	
		SSVS	Common	-48.414	-51.446	-43.495	-4.393	-5.983	2.784	
		SSVS	State-specific	16.783	12.134	10.815	-0.411	2.961	-0.638	
MS-FT THEORETICAL	K = 4	None	Common	7.901	-30.386	4.123	2.898	0.733	-1.624	
		None	State-specific	11.263	12.679	8.562	-0.737	-1.887	-6.043	
		SSVS	Common	-5.469	0.267	2.419	2.278	2.075	-1.958	
		SSVS	State-specific	11.321	13.302	13.524	-0.273	0.219	-1.010	
MS-TVP KITCHEN-SINK	K = 2	SSVS	Common	-39.401	-84.104	-51.234	-4.853	-5.579	-7.278	
		SSVS	State-specific	12.215	9.430	14.090	0.908	0.514	-1.961	
	K = 3	SSVS	Common	-114.149	-156.993	-126.614	-12.394	-10.818	-14.418	
		SSVS	State-specific	23.788	22.774	25.185	-0.577	-0.994	-3.072	
	K = 4	SSVS	Common	-187.724	-190.133	-209.206	-16.165	-17.181	-19.057	
		SSVS	State-specific	21.200	25.655	26.937	-1.168	-1.968	-3.554	
	MS-FT KITCHEN-SINK	K = 2	SSVS	Common	-57.328	-73.540	-42.487	-0.957	-1.397	-1.684
			SSVS	State-specific	12.749	11.300	15.030	-1.861	-1.944	-5.495
K = 3		SSVS	Common	-28.715	-72.196	-64.311	-1.758	-0.587	-8.577	
		SSVS	State-specific	18.781	16.662	17.812	-11.315	-12.182	-14.188	
K = 4		SSVS	Common	-50.802	3.179	5.573	-0.430	-0.716	-2.885	
		SSVS	State-specific	5.515	3.421	3.481	-21.445	-22.107	-23.478	
LINEAR	Taylor rule			2.200	0.357	0.169	-0.601	0.605	-3.410	
	Monetary			-0.324	-1.030	-1.239	-0.875	-0.115	-2.687	
	PPP			0.289	-1.328	-0.407	-0.656	-0.071	-1.127	
	UIP			-0.276	-0.083	-0.190	0.042	-0.087	-0.354	
				SE			UK			
MS-TVP THEORETICAL	K = 4	None	Common	-8.982	-9.025	-2.870	-5.982	-1.321	5.162	
		None	State-specific	-5.147	-0.361	-2.487	-3.696	0.197	3.626	
		SSVS	Common	-6.630	-5.702	1.461	-1.987	1.101	6.036	
		SSVS	State-specific	-5.582	1.893	-0.237	-0.432	1.954	4.312	
MS-FT THEORETICAL	K = 4	None	Common	-1.387	-1.828	-2.744	0.430	2.637	0.912	
		None	State-specific	-0.370	1.454	-2.202	-3.299	-1.581	-0.467	
		SSVS	Common	0.694	1.635	0.271	2.522	3.532	1.887	
		SSVS	State-specific	3.139	3.433	0.806	-0.595	0.148	-0.734	
MS-TVP KITCHEN-SINK	K = 2	SSVS	Common	-15.887	-25.834	-10.620	-1.545	0.370	-0.034	
		SSVS	State-specific	0.412	-0.519	-0.114	-1.090	2.290	2.428	
	K = 3	SSVS	Common	-16.742	-19.861	-13.622	-0.896	1.430	0.443	
		SSVS	State-specific	-1.821	-0.810	-2.402	-1.363	1.245	0.352	
	K = 4	SSVS	Common	-20.343	-24.384	-21.758	-1.955	0.359	-0.078	
		SSVS	State-specific	-2.911	-2.124	-3.380	-2.086	-0.007	-1.260	
	MS-FT KITCHEN-SINK	K = 2	SSVS	Common	-0.319	-1.792	-3.581	0.199	2.924	-1.127
			SSVS	State-specific	-0.941	-3.452	-6.323	-5.674	-4.547	-9.718
K = 3		SSVS	Common	-9.635	-4.012	-3.461	-5.519	-1.261	-3.492	
		SSVS	State-specific	-9.374	-10.639	-14.715	-14.169	-12.835	-14.602	
K = 4		SSVS	Common	-2.274	-2.466	-4.302	-8.277	-5.785	-6.437	
		SSVS	State-specific	-23.645	-25.141	-25.895	-26.066	-25.454	-26.217	
LINEAR	Taylor rule			-4.328	-4.140	-3.620	-4.893	-3.199	-0.890	
	Monetary			-3.632	-4.209	-7.265	-2.725	-1.626	-2.589	
	PPP			-2.508	-2.435	-4.737	-0.652	0.335	-0.748	
	UIP			0.275	-0.025	-0.139	-0.388	-0.373	0.194	

Table 5: Cumulative one-, three-, and twelve-step-ahead LBFs (random walk benchmark) at the end of the full hold-out sample summarized for South Korea, Norway, Sweden and the United Kingdom. Values highlighted green are greater than zero, values highlighted red are smaller than zero, indicating a better or a weaker performance compared to the random walk. Best model in bold.

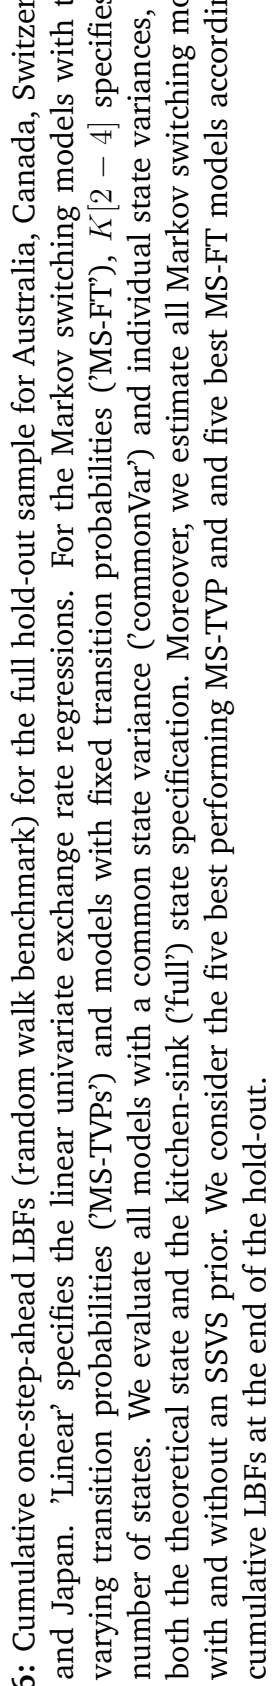


(a) AU



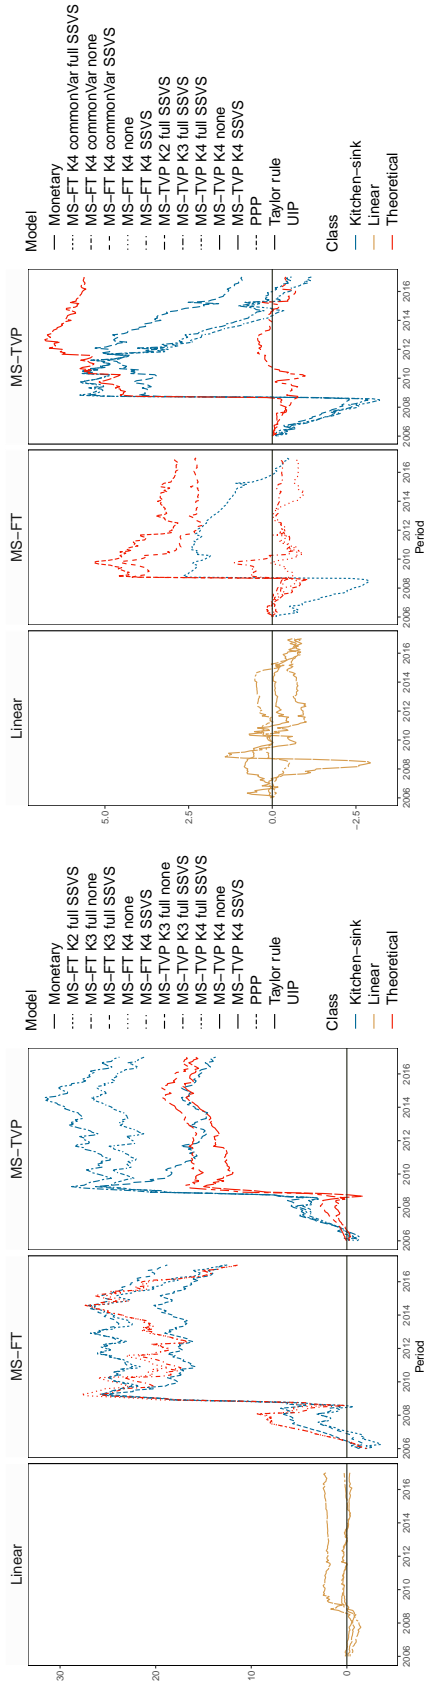
(b) CA

(c) CH

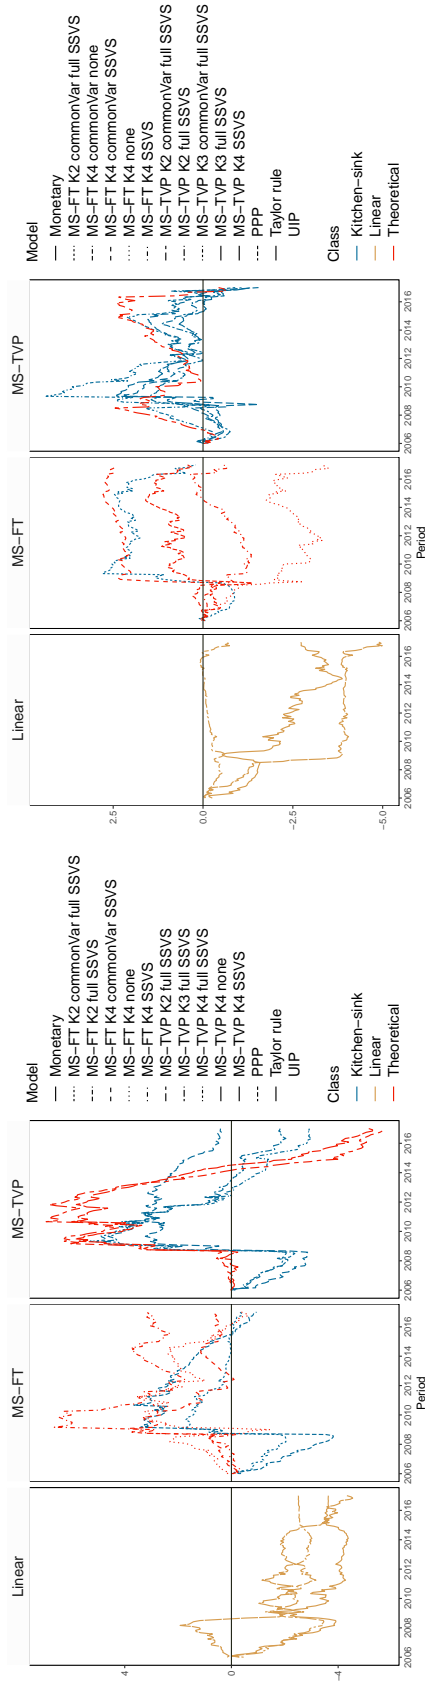


(d) JP

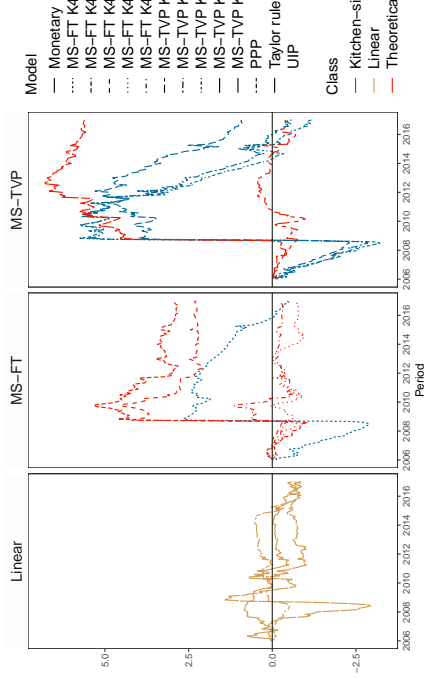
Fig. 6: Cumulative one-step-ahead LBFs (random walk benchmark) for the full hold-out sample for Australia, Canada, Switzerland and Japan. 'Linear' specifies the linear univariate exchange rate regressions. For the Markov switching models with time-varying transition probabilities ('MS-FT') and models with fixed transition probabilities ('MS-FT'), $K[2 - 4]$ specifies the number of states. We evaluate all models with a common state variance ('commonVar') and individual state variances, with both the theoretical state and the kitchen-sink ('full') state specification. Moreover, we estimate all Markov switching models with and without an SSVS prior. We consider the five best performing MS-TVP and five best MS-FT models according to cumulative LBFs at the end of the hold-out.



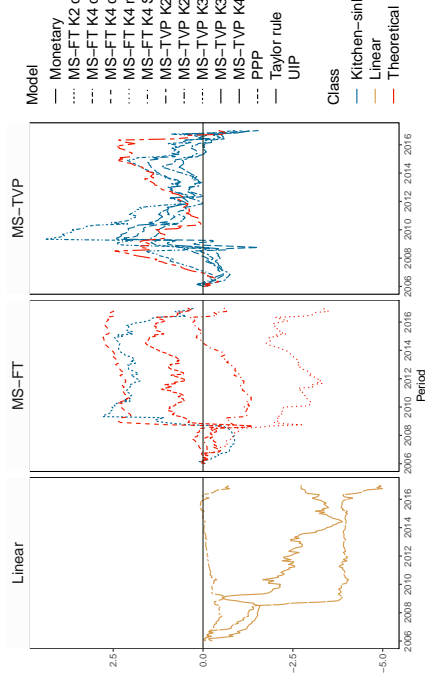
(a) KR



(b) NO



(c) SE



(d) UK

Fig. 7: Cumulative one-step-ahead LBFs (random walk benchmark) for the full hold-out sample for South Korea, Norway, Sweden and the United Kingdom. 'Linear' specifies the linear univariate exchange rate regressions. For the Markov switching models with time-varying transition probabilities ('MS-TVPs') and models with fixed transition probabilities ('MS-FT'), $K[2-4]$ specifies the number of states. We evaluate all models with a common state variance ('commonVar') and individual state variances, with both the theoretical state and the kitchen-sink ('full') state specification. Moreover, we estimate all Markov switching models with and without an SSVS prior. We consider the five best performing MS-TVP and five best MS-FT models cumulative LBFs at the end of the hold-out.

5 Concluding remarks

In this paper, we propose a Bayesian non-linear time series model for exchange rates. Our framework, a multi process model, allows for dynamically switching between selected theoretical exchange rate models that are used to guide the specific choice of covariates included. As an additional novelty, we assume that the transition probabilities vary over time and depend on a measure of the monetary policy stance at home and abroad. This feature enables us to capture breaks in the policy rule of the central bank that, in turn, could impact the prevailing structural exchange rate model adopted. For instance, our framework entails to dynamically switch between models if short-term interest rates hit the zero lower bound.

We use this framework to predict eight exchange rates vis-à-vis the US dollar. Considering the transition probabilities, we find considerable evidence of time-variation. The filtered probabilities indicate that especially after interest rates are bound to zero, model evidence shifts in favor of models other than the Taylor rule based models, highlighting the necessity to control for model uncertainty. To assess whether this feature also translates into predictive accuracy gains, we conduct a forecasting exercise. There, we find that results appear to be rather mixed, with point forecasts being only slightly better than the ones obtained from standard models. In terms of density predictions, however, we observe pronounced accuracy increases for selected exchange rates.

References

- Abbate A and Marcellino M (2018) Point, interval and density forecasts of exchange rates with time varying parameter models. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 181(1), 155–179
- Amisano G and Fagan G (2013) Money growth and inflation: a regime switching approach. *Journal of International Money and Finance* 33, 118–145
- Beckmann J, Koop G, Korobilis D and Schüssler R (2018) Exchange rate predictability and dynamic Bayesian learning. *Essex Finance Centre Working Papers*
- Beckmann J and Schüssler R (2016) Forecasting exchange rates under parameter and model uncertainty. *Journal of International Money and Finance* 60, 267–288
- Belmonte M and Koop G (2014) Model switching and model averaging in time-varying parameter regression models. In *Bayesian Model Comparison*, chapter 2. Emerald Group Publishing Limited, 45–69
- Billio M, Casarin R, Ravazzolo F and Van Dijk HK (2016) Interconnections between eurozone and US booms and busts using a Bayesian panel Markov-switching VAR model. *Journal of Applied Econometrics* 31(7), 1352–1370
- Byrne JP, Korobilis D and Ribeiro PJ (2016) Exchange rate predictability in a changing world. *Journal of International Money and Finance* 62(1), 1–24
- Byrne JP, Korobilis D and Ribeiro PJ (2018) On the sources of uncertainty in exchange rate predictability. *International Economic Review* 59(1), 329–357

- Canova F (1993) Modelling and forecasting exchange rates with a Bayesian time-varying coefficient model. *Journal of Economic Dynamics and Control* 17(1-2), 233–261
- Casarin R, Sartore D and Tronzano M (2018) A Bayesian Markov-switching correlation model for contagion analysis on exchange rate markets. *Journal of Business & Economic Statistics* 36(1), 101–114
- Chinn MD (2006) The (partial) rehabilitation of interest rate parity in the floating rate era: longer horizons, alternative expectations, and emerging markets. *Journal of International Money and Finance* 25(1), 7–21
- Dornbusch R (1976) Expectations and exchange rate dynamics. *Journal of Political Economy* 84(6), 1161–1176
- Engel C (1994) Can the Markov switching model forecast exchange rates? *Journal of International Economics* 36(1-2), 151–165
- Engel C (2014) Exchange rates and interest parity. In *Handbook of international economics*, volume 4. Elsevier, 453–522
- Engel C and West KD (2005) Exchange rates and fundamentals. *Journal of Political Economy* 113(3), 485–517
- Engel C and West KD (2006) Taylor rules and the Deutschmark-Dollar real exchange rate. *Journal of Money, Credit and Banking* 38(5), 1175–1194
- Filardo AJ (1994) Business-cycle phases and their transitional dynamics. *Journal of Business & Economic Statistics* 12(3), 299–308
- Frankel JA (1979) On the mark: a theory of floating exchange rates based on real interest differentials. *The American Economic Review* 69(4), 610–622
- Frühwirth-Schnatter S (2001) Markov chain Monte Carlo estimation of classical and dynamic switching and mixture models. *Journal of the American Statistical Association* 96(453), 194–209
- Frühwirth-Schnatter S (2006) *Finite mixture and Markov switching models*. Springer Science & Business Media
- Frühwirth-Schnatter S and Frühwirth R (2010) Data augmentation and MCMC for binary and multinomial logit models. In *Statistical Modelling and Regression Structures*. Springer, 111–132
- George EI and McCulloch RE (1993) Variable selection via Gibbs sampling. *Journal of the American Statistical Association* 88(423), 881–889
- George EI and McCulloch RE (1997) Approaches for Bayesian variable selection. *Statistica Sinica* 7(2), 339–373
- George EI, Sun D and Ni S (2008) Bayesian stochastic search for VAR model restrictions. *Journal of Econometrics* 142(1), 553–580
- Geweke J and Amisano G (2010) Comparing and evaluating Bayesian predictive distributions of asset returns. *International Journal of Forecasting* 26(2), 216–230
- Groen JJ (2000) The monetary exchange rate model as a long-run phenomenon. *Journal of International Economics* 52(2), 299–319
- Hamilton JD (1994) *Time Series Analysis*. Princeton University Press, Princeton, NJ
- Hooper P and Morton J (1982) Fluctuations in the dollar: a model of nominal and real exchange rate determination. *Journal of International Money and Finance* 1(2), 39–56
- Huber F (2016) Forecasting exchange rates using multivariate threshold models. *The BE Journal of Macroeconomics* 16(1), 193–210

- Huber F (2017) Structural breaks in Taylor rule based exchange rate models - evidence from threshold time varying parameter models. *Economics Letters* 150, 48–52
- Huber F and Fischer MM (2018) A Markov switching factor-augmented VAR model for analyzing US business cycles and monetary policy. *Oxford Bulletin of Economics and Statistics* 80(3), 575–604
- Huber F and Zörner TO (2018) Threshold cointegration in international exchange rates: a Bayesian approach. *International Journal of Forecasting* forthcoming
- Kaufmann S (2015) K-state switching models with time-varying transition distributions—does loan growth signal stronger effects of variables on inflation? *Journal of Econometrics* 187(1), 82–94
- Kim CJ and Nelson CR (1999) State-space models with regime switching: classical and Gibbs-sampling approaches with applications. *MIT Press Books* 1
- Li J, Tsiakas I and Wang W (2015) Predicting exchange rates out of sample: can economic fundamentals beat the random walk? *Journal of Financial Econometrics* 13(2), 293–341
- Mark NC (1995) Exchange rates and fundamentals: evidence on long-horizon predictability. *The American Economic Review* 85(1), 201–218
- Mark NC (2009) Changing monetary policy rules, learning, and real exchange rate dynamics. *Journal of Money, Credit and Banking* 41(6), 1047–1070
- Mark NC and Sul D (2001) Nominal exchange rates and monetary fundamentals: evidence from a small post-Bretton Woods panel. *Journal of International Economics* 53(1), 29–52
- Meese RA and Rogoff K (1983) Empirical exchange rate models of the seventies: do they fit out of sample? *Journal of International Economics* 14(1), 3–24
- Molodtsova T, Nikolsko-Rzhevskyy A and Papell DH (2008) Taylor rules with real-time data: a tale of two countries and one exchange rate. *Journal of Monetary Economics* 55, 63–79
- Molodtsova T, Nikolsko-Rzhevskyy A and Papell DH (2011) Taylor rules and the euro. *Journal of Money, Credit and Banking* 43(2/3), 535–552
- Molodtsova T and Papell DH (2009) Out-of-sample exchange rate predictability with Taylor rule fundamentals. *Journal of International Economics* 77(2), 167–180
- Molodtsova T and Papell DH (2012) Taylor rule exchange rate forecasting during the financial crisis. *NBER International Seminar on Macroeconomics* 9(1), 55–97
- Polson NG, Scott JG and Windle J (2013) Bayesian inference for logistic models using Pólya–Gamma latent variables. *Journal of the American Statistical Association* 108(504), 1339–1349
- Rapach DE and Wohar ME (2002) Testing the monetary model of exchange rate determination: new evidence from a century of data. *Journal of International Economics* 58(2), 359–385
- Sarno L (2005) Towards a solution to the puzzles in exchange rate economics: where do we stand? *Canadian Journal of Economics/Revue canadienne d'économie* 38(3), 673–708
- Sarno L, Valente G and Wohar ME (2004) Monetary fundamentals and exchange rate dynamics under different nominal regimes. *Economic Inquiry* 42(2), 179–193
- Wright JH (2008) Bayesian model averaging and exchange rate forecasts. *Journal of Econometrics* 146(2), 329–341

Appendix A MCMC algorithm

After specifying appropriate starting values, the Gibbs sampler iterates through the following steps:

1. Sample parameters of the measurement equation $\theta = (\beta, \sigma^2, \delta)$ from $p(\theta | \tilde{\mathcal{S}}_T, \Delta \tilde{e}_T)$, with $\Delta \tilde{e}_T = (\Delta e_1, \dots, \Delta e_T)$.
 - (a) Conditional on the exchange rate data \tilde{e}_T and allocation of the states $\tilde{\mathcal{S}}_T$, sampling β and σ^2 can be done in a standard way by drawing the coefficients β_k for $k = 1, \dots, K$ in a block from a multivariate Normal distribution and the variances independently for each state from an inverse Gamma distribution.
 - (b) Conditional on β_k of state $S_t = k$, one is able to sample the elements of δ_k from a Bernoulli conditional posterior distribution.
2. For sampling the unknown states $p(\tilde{\mathcal{S}}_T | \Delta \tilde{e}_T, \theta, \tilde{\mathcal{Z}}_T, \gamma)$ we adapt the filtering algorithm put forth by [Kim and Nelson \(1999\)](#).
3. Following [Polson et al. \(2013\)](#), we sample multinomial coefficients from $p(\gamma | \tilde{\mathcal{S}}_T, \tilde{\mathcal{Z}}_T)$ to construct the time-varying transition probabilities.
4. In case of label switching, we implement an additional permutation step outlined in [Frühwirth-Schnatter \(2001\)](#) ensuring an equal probability of each mode appearing in the posterior distribution.

For the second step, we follow [Kim and Nelson \(1999\)](#) and sample $\tilde{\mathcal{S}}_T$ in a block using a multimove Gibbs sampler. This implies simulating $\tilde{\mathcal{S}}_T$ in block from the following joint conditional distribution

$$p(\tilde{\mathcal{S}}_T | \Delta \tilde{e}_T, \theta, \tilde{\mathcal{Z}}_T, \gamma),$$

and applying a forward filtering and backward smoothing algorithm ([Frühwirth-Schnatter, 2006](#)).

Unfortunately for Step 3, there is no closed form of the multinomial logit posterior distribution. [Frühwirth-Schnatter and Frühwirth \(2010\)](#) therefore introduce two auxiliary layers for estimating state-specific utilities. [Polson et al. \(2013\)](#) take a different approach by representing the likelihood of multinomial logit model as a scale mixture of Gaussians with a Pólya Gamma mixing distribution. In a hierarchical form (by introducing a single layer of latent variables) this strategy implies that the multinomial coefficients are drawn from a set of Gaussians, and the auxiliary variables are sampled

from a Pólya Gamma distribution. This approach has the advantage of fast convergence, simple implementation and no need for an additional layer for approximating the error distribution.

Following Polson et al. (2013), conditional on the latent states $\tilde{\mathbf{S}}_T$ and the auxiliary variables ψ_{jt} for $t = 1, \dots, T$ sampled from a $\mathcal{PG}(\psi_{jt}|1, 0)$ distribution, the posterior quantities are given by

$$\begin{aligned}\bar{\mathbf{V}}_j &= \left(\tilde{\mathbf{Z}}_T' \boldsymbol{\Psi}_j \tilde{\mathbf{Z}}_T + \underline{\mathbf{V}}^{-1} \right)^{-1} \\ \bar{\boldsymbol{\gamma}}_{kj} &= \bar{\mathbf{V}}_j \left(\tilde{\mathbf{Z}}_T' \tilde{\boldsymbol{\kappa}}_j \right)\end{aligned}$$

where $\boldsymbol{\gamma}_{kK}$ is zero for $k = 1, \dots, K$ for reasons of identification, $\boldsymbol{\Psi}_j = \text{diag}(\psi_{j1}, \dots, \psi_{jT})$ and

$$\tilde{\boldsymbol{\kappa}}_j = (\mathcal{I}[S_t = j] - 0.5) - \boldsymbol{\Psi}_j \mathbf{C} \mathbf{e}_j.$$

\mathbf{C} is a $T \times K - 1$ dimensional matrix, with elements being defined as

$$C_{tl} = \log \left(\sum_{k \neq l} \exp(\mathbf{Z}_t' \boldsymbol{\gamma}_{kj}) \right),$$

and \mathbf{e}_j represents the unit vector with an one at the j th position.