Journal of Business & Economic Statistics

ISSN: 0735-0015 (Print) 1537-2707 (Online) Journal homepage: www.tandfonline.com/journals/ubes20

Filtering and Smoothing in State-Space Models with Multiple Regimes*

Nigar Hashimzade, Oleg Kirsanov, Tatiana Kirsanova & Junior Maih

To cite this article: Nigar Hashimzade, Oleg Kirsanov, Tatiana Kirsanova & Junior Maih (27 Apr 2026): Filtering and Smoothing in State-Space Models with Multiple Regimes*, Journal of Business & Economic Statistics, DOI: 10.1080/07350015.2026.2656466

To link to this article: https://doi.org/10.1080/07350015.2026.2656466

Filtering and Smoothing in State-Space Models with Multiple Regimes*

Nigar Hashimzade[a], Oleg Kirsanov[b], Tatiana Kirsanova[b,#], Junior Maih[c]

[a]Department of Economics and Finance, Brunel University of London, United Kingdom

[b]Adam Smith Business School, University of Glasgow, United Kingdom

[c]Norges Bank

[#]trkirsanova@gmail.com

*This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We thank the Editor and two anonymous reviewers for their constructive comments and suggestions, which have substantially improved the paper. All errors remain ours. The authors report there are no competing interests to declare.

Abstract

This paper improves Bayesian filtering techniques in regime-switching state-space models and develops a novel recursion-based smoother for latent variables. The smoother is computationally stable, adaptable to different filters, and easy to implement. We assess its performance in a New Keynesian DSGE model pairing it with three practical filters: the Generalized Pseudo-Bayesian filters of order one (GPB*1*) and two (GPB*2*, often referred to as the Kim or Kim–Nelson filter in applied economics), and the Interacting Multiple Model filter (IMM), common in engineering literature but rarely used in economics. The simulation results show that the IMM filter is about three times faster and at least as accurate as the GPB*2* filter, while our smoother further reduces errors by approximately 25%. Applied to U.S. data from 1947 to 2023, the IMM filter–smoother pair uncovers important monetary policy regime shifts, including those after COVID-19. This demonstrates the practical relevance of the proposed routines for macroeconomic analysis.

Keywords: Markov switching models, Latent variables, Filtering, Smoothing

JEL Reference Number: C11, C32, C54, E52

1 Introduction

Dynamic stochastic general equilibrium (DSGE) models with regime switching have become an important tool for understanding policy changes and structural shifts in macroeconomics. Estimation of such models relies on state–space methods, in which filtering algorithms are used to evaluate the likelihood and reconstruct latent variables, while smoothing algorithms exploit the entire sample to refine estimates of states and regime probabilities.

While filtering methods for regime-switching DSGE models are well established, smoothing remains a weak link. In the macroeconomic literature, the Generalized Pseudo–Bayesian filter of order two (GPB*2*) is commonly implemented following Kim (1994) and Kim and Nelson (1999) and is typically referred to as the Kim–Nelson (KN) filter. [1] Moreover, Kim (1994) also offers a backward-recursion smoother for *regime probabilities*, which is a standard and numerically stable procedure. However, the associated Kim (1994) smoother for *latent variables* has rarely appeared in empirical work. Our own attempts to apply it revealed severe numerical instabilities, which likely explain its absence from the applied literature. Turning to engineering methods, one finds many smoothers for multiple-model estimation of latent variables, but they typically assume an invertible observation-error covariance matrix—a natural assumption in engineering but incompatible with standard DSGE practice, where measurement errors are often set to zero. This mismatch leaves researchers without a reliable latent-variable smoothing tool for regime-switching macro models.

The main contribution of this paper is to fill this gap. We develop a new recursion-based smoother, building on the disturbance-smoothing approach of De Jong (1988) and Durbin and Koopman (2012), and adapt it to Markov-switching state–space models in a way that works robustly even when measurement errors are absent. To our knowledge, this is the first smoothing algorithm tailored to the needs of regime-switching DSGE analysis.

Because smoothers must be paired with filters, we also revisit the filtering side. In practice, macroeconomists have relied almost exclusively on the KN filter. By contrast, the engineering literature overwhelmingly favors the Interactive Multiple Model (IMM) filter, valued for its speed–accuracy trade-off (Blom, 1984; Blom and Bar-Shalom, 1988). We introduce the IMM filter into the macroeconomic context and study it alongside the GPB family, including the KN filter, showing in our experiments that IMM matches the KN accuracy while running about three times faster.

The contributions of the paper are therefore threefold. First, we develop a novel smoother tailored to regime-switching DSGE models with zero measurement errors, and we demonstrate its properties through extensive Monte Carlo analysis of a medium-scale New Keynesian model. The new smoother reduces estimation errors of latent variables by about 25% for sample sizes of 100–300 quarters, a range typical in applied macroeconomic work. We further validate its performance by benchmarking against particle-based methods with well-defined asymptotic properties, introducing a small observation error where required to ensure comparability of likelihoods. Second, we clarify the relationship between IMM and GPB filters, situating our

work within both the econometrics and engineering literature. Third, we illustrate the practical usefulness of our methodology with an application to U.S. data from 1947 to 2023, highlighting shifts in monetary policy regimes, including the post-COVID-19 period. To ensure accessibility and reproducibility, we make all methods available in the RISE toolbox [2] for MATLAB.

The remainder of the paper is organized as follows. Section 2 reviews the filtering framework, including the GPB and IMM families. Section 3 presents our new recursion-based smoother. Section 4 reports the simulation evidence. Section 5 provides an empirical application to U.S. data. Section 6 concludes.

2 Multiple-Regime Filters

2.1 The Filtering Problem

We start with a general multiple-regime state-space representation of a linear discrete-time dynamic model consisting of a measurement equation (1) and a transition equation (2),

$$y_t = c_{y,s_t} + Z_{s_t}\alpha_t + g_{s_t}\varepsilon_t, \quad (1)$$

$$\alpha_t = c_{\alpha,s_t} + T_{s_t}\alpha_{t-1} + R_{s_t}v_t, \quad (2)$$

where y_t is a $p\times 1$ vector of observations, α_t is an $m\times 1$ vector of unobserved state variables, and ε_t and v_t are independent standard Gaussian random variables, $t=1,\ldots,n$. All model parameters $\left\{c_{y,s_t}, c_{\alpha,s_t}; Z_{s_t}, T_{s_t}; g_{s_t}, R_{s_t}\right\}$ depend on regime s_t, which is an outcome of a Markov process with $h\geq 1$ discrete regimes. This process is described by the transition probability matrix with the generic element $Q(s_{t-1},s_t)=\Pr\left[s_t \mid s_{t-1}\right]$, so that $\sum_{s_t=1}^{h} Q(s_{t-1},s_t)=1$ for every regime s_{t-1} and every time *t*.

The information available at time *t* is fully contained in the vector of observations $Y_t := \{y_1,\ldots,y_t\}$. The object of interest is an estimate of the unobserved state vector α_t, for which three estimators, $\alpha_{t|t-1}$, $\alpha_{t|t}$ and $\alpha_{t|h}$, are available in the Bayesian framework. The first estimator is the forecast of α_t based on information Y_{t-1},

$$\alpha_{t|t-1} := \mathbb{E}\left[\alpha_t \mid Y_{t-1}\right].$$

Its mean square error (MSE) is defined as

$$P_{t|t-1} := \mathbb{E}\left[\left(\alpha_t-\alpha_{t|t-1}\right)\left(\alpha_t-\alpha_{t|t-1}\right)' \mid Y_{t-1}\right].$$

In the single-regime linear setting with Gaussian shocks, these objects and the associated likelihood $f\left(y_t \mid Y_{t-1}\right)$ are computed by the well-established technique of the standard Kalman filter (KF), which in this case is exact and optimal (Kalman, 1960). Working in a multiple-regime environment is more challenging because of the explosive dimensionality of the problem.

Specifically, in a multiple-regime environment, exact estimation is infeasible because the number of histories that a Kalman-type filter needs to take into account increases exponentially with every time period. At any given time *t*, a multiple-regime dynamic system can be in one of *h* possible regimes, each corresponding to a realization of *h* mutually exclusive and exhaustive random events. Denote the sequence of realized regimes from the beginning of observations up to time *t* by $\mathcal{H}_t$:

$$\mathcal{H}_t = \{s_1, s_2, ..., s_{t-1}, s_t\} \in \mathbb{H}_{t,t}$$

where $\mathbb{H}_{N,t}$ is the set of all possible histories of length *N* that end in the period *t*. There are h^t possible mutually exclusive and exhaustive histories up to time *t*. Let $\mathcal{H}_t^h$ denote one possible history, $\mathcal{H}_t^h = \{s_1^h, s_2^h, ..., s_{t-1}^h, s_t^h\}$. Using the total probability theorem, the conditional pdf at time *t* is obtained as a Gaussian mixture with the number of terms equal to h^t:

$$f\left(y_{t+1} \mid Y_t\right) = \sum_{h=1}^{h^t} f\left(y_{t+1} \mid \mathcal{H}_t = \mathcal{H}_t^h, Y_t\right) \Pr\left[\mathcal{H}_t = \mathcal{H}_t^h \mid Y_t\right].$$

The probability of a given regime history is computed using the Bayes formula:

$$\begin{aligned} \Pr\left[\mathcal{H}_t = \mathcal{H}_t^h \mid Y_t\right] &= \frac{f\left(y_t \mid \mathcal{H}_t = \mathcal{H}_t^h, Y_{t-1}\right) \Pr\left[s_t = s_t^h \mid \mathcal{H}_{t-1} = \mathcal{H}_{t-1}^h, Y_{t-1}\right]}{f\left(y_t \mid Y_{t-1}\right)} \\ &\times \Pr\left[\mathcal{H}_{t-1} = \mathcal{H}_{t-1}^h \mid Y_{t-1}\right] \end{aligned}$$

When the regime switches have the Markov property, the second term in the numerator simplifies to $\Pr\left[s_t \mid \mathcal{H}_{t-1} = \mathcal{H}_{t-1}^h, Y_{t-1}\right] = \Pr\left[s_t \mid s_{t-1}\right] = Q(s_{t-1}, s_t)$. However, the knowledge of the entire past history is still needed for the terms $f\left(y_t \mid \mathcal{H}_t = \mathcal{H}_t^h, Y_{t-1}\right)$ and $\Pr\left[\mathcal{H}_{t-1} = \mathcal{H}_{t-1}^h \mid Y_{t-1}\right]$, even if the regimes follow a Markov process. This is known as the path dependence problem.

In practice, one has to resort to some approximation. In this sense, all practical multiple-regime filters are approximate and, therefore, suboptimal. [3] We focus on two practical families of approximate filters that address path dependence in different ways: the Generalized Pseudo–Bayesian (GPB) filters and the Interacting Multiple Model (IMM) filters.

The indexing of GPB filters is based on the length of tracked history (sometimes referred to as "depth"). Following common notation, the GPB*N* filter uses information from the previous *N*

periods, including the current one, and so tracks h^N histories. Thus, GPB*1* ignores past history and uses only the current period, while GPB*2* incorporates information from the current period and one preceding period, maintaining h^2 possible histories when moving to the next step, and so on. In the economics literature, GPB*2* is commonly referred to as the Kim–Nelson (KN) filter, following Kim (1994) and Kim and Nelson (1999). To avoid ambiguity, in what follows we will use the notation GPB*2*, reserving "KN" only for historical references.

For the IMM family, we first present the IMM algorithm as developed in Blom (1984) and Blom and Bar-Shalom (1988). It also uses information from the current period and one immediately preceding period, but, unlike the GPB*2*, it combines h^2 possible histories at every step into *h* and, thus, maintains only h^1 histories when moving to the next step. Therefore, the commonly used IMM is, technically, IMM*1*. In principle, IMM*N* can be constructed and employed. However, Blom and Bar-Shalom (1988) suggest that higher-order GPB and IMM are related. We prove that higher-order GPB and IMM are equivalent under matching initialization. This equivalence may explain scarce mention of higher-order IMM in the literature.

While we use the general notations in the exposition in Sections 2 and 3, the description of the main filter and smoother algorithms in the form that enable easy coding is given in Appendix A.

2.2 Two Practical Families of Filters

The exposition of commonly used filters serves a dual purpose. First, it introduces the notations used in the description of the smoother in the next section. Second, it will help to show, in the subsequent analysis of the IMM family, the equivalence of higher-order IMM and GPB filters. To facilitate comparison, before proceeding to the multiple-regime filters we state the standard (single-regime) KF, associated with system (1)-(2) omitting s_t, as the following set of equations:

Forecast

$$\alpha_{t|t-1} = c_{\alpha,t} + T_t \alpha_{t-1|t-1},$$
$$P_{t|t-1} = T_t P_{t-1|t-1} T_t' + R_t R_t',$$

Update

$$\alpha_{t|t} = \alpha_{t|t-1} + P_{t|t-1} Z_t' \left[F_{t|t-1}\right]^{-1} \nu_{t|t-1},$$
$$P_{t|t} = \left(I - P_{t|t-1} Z_t' \left[F_{t|t-1}\right]^{-1} Z_t\right) P_{t|t-1},$$

where

$$\nu_{t|t-1} = y_t - Z_t \alpha_{t|t-1} - c_{y,t}, \quad F_{t|t-1} = Z_t P_{t|t-1} Z_t' + H_t, \quad H_t = g_t g_t',$$

and the associated likelihood is

$$\Lambda_t = f\left(y_t \mid Y_{t-1}\right) = (2\pi)^{-p/2} \left|F_{t|t-1}\right|^{-1/2} \exp\left(-\frac{1}{2} \nu_{t|t-1}' \left[F_{t|t-1}\right]^{-1} \nu_{t|t-1}\right)$$

In what follows, we will describe the mapping from the inputs to the outputs of the KF as a multidimensional function $\mathcal{K}(\cdot)$:

$$\left(\alpha_{t|t-1}, P_{t|t-1}, \alpha_{t|t}, P_{t|t}; \Lambda_t\right) = \mathcal{K}\left(c_{\alpha,t}, T_t, R_t; c_{y,t}, Z_t, g_t; \alpha_{t-1|t-1}, P_{t-1|t-1}; y_t\right).$$

We also use the following notations.

Let $\mathcal{H}_t$ denote the history of regimes in *N* consecutive periods ending with period *t*,

$$\mathcal{H}_t := \{s_{t-N+1}, ..., s_{t-1}, s_t\} \in \mathbb{H}_{N,t},$$

and let $\mathcal{C}_t$ be the "collapsed" history, defined as

$$\mathcal{C}_t := \{s_{t-N+2}, ..., s_t\} \in \mathbb{H}_{N-1,t}.$$

Hence $\mathcal{H}_t = \{\mathcal{C}_{t-1}, s_t\} = \{s_{t-N+1}, \mathcal{C}_t\}$, and $\mathcal{H}_{t-1} = \{s_{t-N}, ..., s_{t-2}, s_{t-1}\} \in \mathbb{H}_{N,t-1}$.

Let

$$\mu_{t|t}^{(\mathcal{H}_t)} := \Pr\left[\mathcal{H}_t \mid Y_t\right]$$

be the probability of the realization of a particular history $\mathcal{H}_t$ conditional on information at time *t*.

2.2.1 Family of GPB Filters

The GPB algorithm of order *N*, denoted GPB*N*, takes into account all h^N possible histories of the fixed length *N*, finishing at the current time period. It is implemented as follows.

Define

$$\mu_{t|t}^{(\mathcal{C}_t)} := \Pr\left[\mathcal{C}_t \mid Y_t\right] = \sum_{s_{t-N+1}=1}^{h} \mu_{t|t}^{(\mathcal{H}_t)}$$

as the probability of collapsed history, $\mathcal{C}_t$, conditional on information at time *t*.

Algorithm 1. *GPBN Algorithm*

Step 0. Start at $t = 1$. Initialize h^{N-1} versions of the state vector $\alpha_{t-1|t-1}^{(\mathcal{C}_{t-1})}$, the MSE matrix $P_{t-1|t-1}^{(\mathcal{C}_{t-1})}$, and probabilities $\mu_{t-1|t-1}^{(\mathcal{C}_{t-1})}$, $\mathcal{C}_{t-1} \in \mathbb{H}_{N-1,t-1}$.

Step 1. Compute h^N standard KF forecasts to obtain

$$\left(\alpha_{t|t-1}^{(\mathcal{H}_t)}, P_{t|t-1}^{(\mathcal{H}_t)}, \alpha_{t|t}^{(\mathcal{H}_t)}, P_{t|t}^{(\mathcal{H}_t)}; \Lambda_t^{(\mathcal{H}_t)}\right) = \mathcal{K}\left(c_{\alpha,s_t}, T_{s_t}, R_{s_t}; c_{y,s_t}, Z_{s_t}, g_{s_t}; \alpha_{t-1|t-1}^{(\mathcal{C}_{t-1})}, P_{t-1|t-1}^{(\mathcal{C}_{t-1})}; y_t\right)$$

Step 2. Compute the updated probabilities $\mu_{t|t}^{(\mathcal{H}_t)}$ according to

$$\begin{aligned}
&\mu_{t|t}^{(\mathcal{H}_t)} = \Pr\left[\mathcal{H}_t |\ Y_t\right] = \Pr\left[\mathcal{H}_t |\ Y_{t-1}, y_t\right] \\
&= \frac{f\left(y_t |\ \mathcal{H}_t, Y_{t-1}\right)\Pr\left[s_t |\ \mathcal{C}_{t-1}, Y_{t-1}\right]}{f\left(y_t |\ Y_{t-1}\right)}\Pr\left[\mathcal{C}_{t-1} |\ Y_{t-1}\right] \\
&= \frac{\Lambda_t^{(\mathcal{H}_t)} Q\left(s_{t-1}, s_t\right)}{\sum_{\mathcal{H}_t}\Lambda_t^{(\mathcal{H}_t)} Q\left(s_{t-1}, s_t\right)\mu_{t-1|t-1}^{(\mathcal{C}_{t-1})}}\mu_{t-1|t-1}^{(\mathcal{C}_{t-1})},
\end{aligned}$$

where in the last line we used the Markov property.

$$\Pr\left[s_t |\ \mathcal{C}_{t-1}, Y_{t-1}\right] = \Pr\left[s_t |\ s_{t-1}\right] = Q\left(s_{t-1}, s_t\right), \quad (4)$$

and

$$\begin{aligned}
f\left(y_t |\ Y_{t-1}\right) &= \sum_{\mathcal{H}_t} f\left(y_t |\ Y_{t-1}, \mathcal{H}_t\right)\Pr\left[\mathcal{H}_t |\ Y_{t-1}\right] \quad (5) \\
&= \sum_{\mathcal{H}_t} f\left(y_t |\ Y_{t-1}, \mathcal{H}_t\right)\Pr\left[s_t |\ \mathcal{C}_{t-1}, Y_{t-1}\right]\Pr\left[\mathcal{C}_{t-1} |\ Y_{t-1}\right] \\
&= \sum_{\mathcal{H}_t}\Lambda_t^{(\mathcal{H}_t)} Q\left(s_{t-1}, s_t\right)\mu_{t-1|t-1}^{(\mathcal{C}_{t-1})},
\end{aligned}$$

where the sum is taken over all possible realizations of $\mathcal{H}_t$.

Step 3. Next, h^N KF outputs are merged into h^{N-1} using conditional probabilities $\Pr[s_{t-N+1} | Y_t, \mathcal{C}_t]$ computed as:

$$\Pr[s_{t-N+1} | Y_t, \mathcal{C}_t] = \frac{\Pr\left(s_{t-N+1}, \mathcal{C}_t | Y_t\right)}{\Pr\left(\mathcal{C}_t | Y_t\right)} = \frac{\Pr\left(\mathcal{H}_t | Y_t\right)}{\Pr\left(\mathcal{C}_t | Y_t\right)} = \frac{\mu_{t|t}^{(\mathcal{H}_t)}}{\sum_{s_{t+1-N}=1}^{h}\mu_{t|t}^{(\mathcal{H}_t)}}.$$

Thus,

$$\alpha_{t|t}^{(\mathcal{C}_t)} = \sum_{s_{t-N+1}=1}^{h}\Pr[s_{t-N+1} | Y_t, \mathcal{C}_t]\alpha_{t|t}^{(\mathcal{H}_t)}, \quad (6)$$

$$P_{t|t}^{(\mathcal{C}_t)} = \sum_{s_{t-N+1}=1}^{h} \Pr[s_{t-N+1} \mid Y_t, \mathcal{C}_t] \left\{ P_{t|t}^{(\mathcal{H}_t)} + \left(\alpha_{t|t}^{(\mathcal{C}_t)} - \alpha_{t|t}^{(\mathcal{H}_t)} \right) \left(\alpha_{t|t}^{(\mathcal{C}_t)} - \alpha_{t|t}^{(\mathcal{H}_t)} \right)' \right\}, \quad (7)$$

$$\mu_{t|t}^{(\mathcal{C}_t)} = \Pr\left[\mathcal{C}_t \mid Y_t \right] = \sum_{s_{t-N+1}=1}^{h} \Pr\left[\mathcal{H}_t \mid Y_t \right] = \sum_{s_{t-N+1}=1}^{h} \mu_{t|t}^{(\mathcal{H}_t)}. \quad (8)$$

Go to Step 1. Updated $\alpha_{t|t}^{(\mathcal{C}_t)}$, $P_{t|t}^{(\mathcal{C}_t)}$ and $\mu_{t|t}^{(\mathcal{C}_t)}$ serve as initializations for the next time period ($t = 2, 3, ..., n$) in a recursion at Step 1.

The *t*-increment likelihood $L_t = \log f\left(y_t \mid Y_{t-1} \right),$ is computed using (5) as part of the filter algorithm.

To continue the recursion to the end of the sample we only need to compute $\left\{ \alpha_{t|t}^{(\mathcal{C}_t)}, P_{t|t}^{(\mathcal{C}_t)}, \mu_{t|t}^{(\mathcal{C}_t)} \right\}_{(\mathcal{C}_t) \in \mathbb{H}_{N-1,t}}$ at each step of the algorithm. These quantities are also used to compute the (averaged over regime sequences) state vectors and the MSE matrices $\left\{ \alpha_{t|t}, P_{t|t} \right\}_{t=1:n}$ for each time *t*, and the probability $\mu_t^{(s_t)}$ of the system being in regime s_t conditional on information at time *t*:

$$\alpha_{t|t} = \sum_{\mathcal{C}_t} \Pr\left[\mathcal{C}_t \mid Y_t \right] \alpha_{t|t}^{(\mathcal{C}_t)} = \sum_{\mathcal{C}_t} \mu_{t|t}^{(\mathcal{C}_t)} \alpha_{t|t}^{(\mathcal{C}_t)}, \quad (9)$$

$$P_{t|t} = \sum_{\mathcal{C}_t} \Pr\left[\mathcal{C}_t \mid Y_t \right] \left\{ P_{t|t}^{(\mathcal{C}_t)} + \left(\alpha_{t|t} - \alpha_{t|t}^{(\mathcal{C}_t)} \right) \left(\alpha_{t|t} - \alpha_{t|t}^{(\mathcal{C}_t)} \right)' \right\}, \quad (10)$$

$$\begin{aligned} &\sum_{\mathcal{C}_t} \mu_{t|t}^{(\mathcal{C}_t)} \left\{ P_{t|t}^{(\mathcal{C}_t)} + \left(\alpha_{t|t} - \alpha_{t|t}^{(\mathcal{C}_t)} \right) \left(\alpha_{t|t} - \alpha_{t|t}^{(\mathcal{C}_t)} \right)' \right\}, \\ &\mu_{t|t}^{(s_t)} = \Pr[s_t \mid Y_t] = \sum_{\mathcal{C}_{t-1}} \Pr[\mathcal{H}_t \mid Y_t] = \sum_{\mathcal{C}_{t-1}} \mu_{t|t}^{(\mathcal{H}_t)}. \end{aligned} \quad (11)$$

These objects can be computed outside the recursion.

This algorithm works for any $N \geq 1$. Note that for $N = 1$ we have $\mathcal{C}_t = \varnothing, \mathcal{H}_t = s_t$. This implies that the GPB*1* algorithm considers possible regime histories only at their latest instant and merges all preceding regime sequences into one, using *common* initial conditions $\alpha_{t-1|t-1}$ at each time step.

2.2.2 Family of IMM Filters

The IMM filter maintains the dimensionality of the histories at *h* in each time period. This is achieved by combining, or mixing, histories in every recursion of the algorithm immediately after making a step forward in time, which would otherwise result in an increase of dimensionality to h^2. Effectively, at every *t*, mixing replaces a set of h^2 "exact"histories by a reduced set of *h* approximate histories weighted by the probabilities of transition from the previous state, s_{t-1}, to s_t. These *h* histories are then filtered and updated in the usual way.

We outline the IMM algorithm—denoted IMM*1*—as it only carries h^1 histories—to enable its direct comparison with GPB*1* and GPB*2* algorithms, in order to pinpoint the dimension reduction [4] at the different stages of the algorithms. In the IMM*1*, histories are combined, or mixed, after the state update and before the measurement update, and in the GPB*2*, histories are combined, or collapsed after the state and measurement updates. [5]

Algorithm 2. *IMM1 Algorithm.*

Step 0. Initialize *h* versions of the state vector $\alpha_{t-1|t-1}^{(s_{t-1})}$, the MSE matrix $P_{t-1|t-1}^{(s_{t-1})}$, and regime probabilities $\mu_{t-1|t-1}^{(s_{t-1})}$, $s_{t-1}=1,...,h.$

Step 1. Compute the *mixing probabilities* defined as

$$\mu_{t-1|t-1}^{(s_{t-1}|s_t)} := \Pr\left[s_{t-1} \mid Y_{t-1}, s_t \right].$$

Note that

$$\begin{aligned}\Pr[s_{t-1}, s_t \mid Y_{t-1}] &= \Pr\left[s_{t-1} \mid Y_{t-1}, s_t \right] \Pr[s_t \mid Y_{t-1}] \\ &= \Pr\left[s_t \mid s_{t-1}, Y_{t-1} \right] \Pr[s_{t-1} \mid Y_{t-1}].\end{aligned}$$

Therefore,

$$\begin{aligned}\mu_{t-1|t-1}^{(s_{t-1}|s_t)} &= \Pr\left[s_{t-1} \mid Y_{t-1}, s_t \right] \\ &= \frac{\Pr\left[s_t \mid s_{t-1}, Y_{t-1} \right] \Pr[s_{t-1} \mid Y_{t-1}]}{\sum_{s_{t-1}} \left(\Pr\left[s_t \mid s_{t-1}, Y_{t-1} \right] \Pr[s_{t-1} \mid Y_{t-1}] \right)} \\ &= \frac{Q(s_{t-1}, s_t)\mu_{t-1|t-1}^{(s_{t-1})}}{\sum_{s_{t-1}} Q(s_{t-1}, s_t)\mu_{t-1|t-1}^{(s_{t-1})}}.\end{aligned}$$

Step 2. Compute the *mixed* state vectors and MSE matrices for each regime:

$$\hat{\alpha}_{t-1|t-1}^{(*,s_t)} = \sum_{s_{t-1}} \Pr\left[s_{t-1} \mid Y_{t-1}, s_t \right] \alpha_{t-1|t-1}^{(s_{t-1})} = \sum_{s_{t-1}} \mu_{t-1|t-1}^{(s_{t-1}|s_t)} \alpha_{t-1|t-1}^{(s_{t-1})}, \quad (13)$$

$$\hat{P}_{t-1|t-1}^{(*,s_t)} = \sum_{s_{t-1}} \Pr\left[s_{t-1} \mid Y_{t-1}, s_t\right] \left\{ P_{t-1|t-1}^{(s_{t-1})} + \left(\alpha_{t-1|t-1}^{(s_{t-1})} - \hat{\alpha}_{t-1|t-1}^{(*,s_t)}\right)\left(\alpha_{t-1|t-1}^{(s_{t-1})} - \hat{\alpha}_{t-1|t-1}^{(*,s_t)}\right)' \right\} \quad (14)$$

$$= \sum_{s_{t-1}}^{h} \mu_{t-1|t-1}^{(s_{t-1}, s_t)} \left\{ P_{t-1|t-1}^{(s_{t-1})} + \left(\alpha_{t-1|t-1}^{(s_{t-1})} - \hat{\alpha}_{t-1|t-1}^{(*,s_t)}\right)\left(\alpha_{t-1|t-1}^{(s_{t-1})} - \hat{\alpha}_{t-1|t-1}^{(*,s_t)}\right)' \right\}.$$

Once the states and MSEs are mixed, the memory of s_{t-1} is "cleared", so we put an asterisk in place of the disappeared index s_{t-1}. This reduces the dimensionality from h^2 to h.

Step 3. For each history compute the standard KF to obtain

$$\left(\alpha_{t|t-1}^{(s_t)}, P_{t|t-1}^{(s_t)}, \alpha_{t|t}^{(s_t)}, P_{t|t}^{(s_t)}; \Lambda_t^{(s_t)}\right) = \mathcal{K}\left(c_{\alpha,s_t}, T_{s_t}, R_{s_t}; c_{y,s_t}, Z_{s_t}, g_{s_t}; \hat{\alpha}_{t-1|t-1}^{(*,s_t)}, \hat{P}_{t-1|t-1}^{(*,s_t)}; y_t\right)$$

Note that the KF inputs are contemporary: as the memory s_{t-1} is mixed out, the timing of system matrices and states coincide at s_t.

Step 4. Update the probabilities $\mu_{t|t}^{(s_t)}$,

$$\begin{aligned}
\mu_{t|t}^{(s_t)} &= \Pr\left[s_t \mid Y_t\right] = \Pr\left[s_t \mid Y_{t-1}, y_t\right] \\
&= \frac{f\left(y_t \mid s_t, Y_{t-1}\right) \sum_{s_{t-1}} \Pr\left[s_t \mid s_{t-1}, Y_{t-1}\right] \Pr\left[s_{t-1} \mid Y_{t-1}\right]}{f\left(y_t \mid Y_{t-1}\right)} \quad (15) \\
&= \frac{\Lambda_t^{(s_t)} \sum_{s_{t-1}} Q(s_{t-1}, s_t) \mu_{t-1|t-1}^{(s_{t-1})}}{\sum_{s_t} \Lambda_t^{(s_t)} \sum_{s_{t-1}} Q(s_{t-1}, s_t) \mu_{t-1|t-1}^{(s_{t-1})}},
\end{aligned}$$

where in the last line we used

$$\begin{aligned}
f\left(y_t \mid Y_{t-1}\right) &= \sum_{s_t} f\left(y_t \mid Y_{t-1}, s_t\right) \Pr\left[s_t \mid Y_{t-1}\right] \quad (16) \\
&= \sum_{s_t} f\left(y_t \mid Y_{t-1}, s_t\right) \sum_{s_{t-1}} \left(\Pr\left[s_t \mid s_{t-1}, Y_{t-1}\right] \Pr\left[s_{t-1} \mid Y_{t-1}\right] \right) \\
&= \sum_{s_t} \Lambda_t^{(s_t)} \sum_{s_{t-1}} Q(s_{t-1}, s_t) \mu_{t-1|t-1}^{(s_{t-1})}
\end{aligned}$$

Go to Step 1. The KF outputs and the updated probabilities $\left\{\alpha_{t|t}^{(s_t)}, P_{t|t}^{(s_t)}, \mu_{t|t}^{(s_t)}\right\}$ serve as initializations for the next time step in a recursion.

The *t*-increment likelihood $L_t = \log f\left(y_t \mid Y_{t-1}\right)$ is computed using formula (16) as part of the filtering algorithm. The state vectors and MSE matrices $\left\{\alpha_{t|t}, P_{t|t}\right\}_{t=1:n}$ for each time *t* are computed using the same formulae (9)-(10) as in the GPB algorithm.

A comparison of IMM*1* to GPB*1* [6] shows that the essential difference is in the objects used as inputs in the KF. In GPB*1* these are the fused past states, which are independent of the regime, as GPB*1* keeps no past history at all. In IMM*1* the past states (one period back) are mixed out before the state vectors and the MSEs are fed into the KF. In other words, GPB*1* mixes *h* possible current states, s_t, with probabilities $\mu_{t|t}^{(s_t)}$, while IMM*1* mixes *h* possible previous states, s_{t-1}, for each of the *h* possible current states, s_t, with conditional probabilities $\mu_{t-1|t-1}^{(s_{t-1}|s_t)}$.

2.2.3 Extending IMM*1* to a Higher Order: The Equivalence Result

Consider the problem of extending IMM*1* to IMM*2*, using the same idea that for dimension reduction, some earlier states should be mixed out.

We now have three state indices, $\left\{s_{t-2}, s_{t-1}, s_t\right\}$, and we need to initialize the filter with h^2 values of the state vector $\alpha_{t-1|t-1}^{(s_{t-2},s_{t-1})}$, the MSE matrix $P_{t-1|t-1}^{(s_{t-2},s_{t-1})}$, and regime probabilities $\mu_{t-1|t-1}^{(s_{t-2},s_{t-1})}$. There are two options for dimension reduction. We can either mix out the earliest tracked state only, s_{t-2}, or we can mix out two previous tracked states, $\left\{s_{t-2}, s_{t-1}\right\}$. There is no obvious reason for preferring one option to another. In fact, it is straightforward to show that the first option leads to the GPB*2* and the second option leads to the IMM*1*. This equivalence has been noted informally in the engineering literature (see Blom and Bar-Shalom, 1988); here we provide a formal statement and its proof. The result holds in general for any $N \geq 2$, as stated in the theorem below.

Theorem 1 (IMM-GPB equivalence).

For any $N \geq 2$ the IMMN filtering algorithm is reducible to IMM1 and to any of the GPBK for $2 \leq K \leq N$, by selecting the length of past history, $\left\{s_{t-N}, \ldots, s_{t-K}\right\}$, to be mixed out for dimension reduction.

The proof is presented in Appendix B. This proof includes the exact initialization conditions that ensure identical outcomes of the IMM and GPB filters. Because of this equivalence, in the rest of the paper we only investigate IMM*1* and we omit the order index where practical.

3 Smoothing of State Vectors

After estimating the states and regime probabilities through forward propagation, we can improve the inference on s_t and α_t using the information from the entire sample. This process, known as smoothing, is usually conducted by backward propagation.

A practical algorithm for computing smoothed *regime probabilities* is detailed in Kim (1994). In this paper, we generalize this probability recursion to arbitrary tracked history so it is applicable to higher-order filters; see Appendix C.

However, it is well known that smoothing *state vectors*, even in a single-regime model, can be challenging, and is even more so in models with multiple regimes. Kim (1994) introduces an algorithm specifically designed to work with the GPB*2* filter. Unfortunately, Kim's smoother is computationally unstable as it requires inverting large auxiliary matrices, that do not appear at the filtering stage. Existing smoothers developed for engineering applications typically exploit measurement errors and require the invertibility of matrix $H_{s_t} = g_{s_t} g_{s_t}^{'}$ [7], a condition that is often not met in economic applications. In this paper, we construct a novel multiple-regime smoother, drawing on a single-regime smoothing algorithm in Durbin and Koopman (2012) and De Jong (1988). This smoother is based on a recursive algorithm that does not rely on invertibility of observation-error covariance matrices, and only requires matrix inversions that are part of the corresponding filter and would have been computed at the filtering stage. It can be applied to a model with Markov switching in conjunction with a filter with arbitrary length of tracked history.

In principle, filtering provides real–time estimates of the states using information up to *t*, while smoothing revises these estimates in light of information that arrives later. Running backward from the end of the sample, each state estimate is adjusted according to how well it predicted the next observation: if the forecast error at $t+1$ was positive, the smoother nudges the estimate of latent variable α_t upwards; if negative, it nudges α_t down. This backward correction reduces variance and corrects distortions caused by initialization. Existing single-regime algorithms such as those of De Jong (1988) and Durbin and Koopman (2012) implement this logic effectively by reusing the filter's prediction errors and gains, avoiding the need to form or invert large joint covariance matrices. These single–regime smoothers provide the foundation for the recursion we adapt to the regime–switching case.

By definition, smoothed state vectors and MSE matrices are:

$$\alpha_{t|h}^{(\mathcal{H}_t)} := \mathbb{E}\left[\alpha_t^{(\mathcal{H}_t)} \mid Y_n, \mathcal{H}_t \right], \quad (17)$$

$$P_{t|h}^{(\mathcal{H}_t)} := \mathbb{E}\left[\left(\alpha_t^{(\mathcal{H}_t)} - \alpha_{t|t-1}^{(\mathcal{H}_t)} \right) \left(\alpha_t^{(\mathcal{H}_t)} - \alpha_{t|t-1}^{(\mathcal{H}_t)} \right)^{'} \mid Y_n, \mathcal{H}_t \right]. \quad (18)$$

Define the forecast error of the state vector at time *t* with history s_t as

$$\xi_{t|t-1}^{(\mathcal{H}_t)} := \alpha_t - \alpha_{t|t-1}^{(\mathcal{H}_t)}.$$

Then

$$P_{t|t-1}^{(\mathcal{H}_t)} = \mathbb{E}\left[\xi_{t|t-1}^{(\mathcal{H}_t)} \xi_{t|t-1}^{(\mathcal{H}_t)'} \mid Y_{t-1}, \mathcal{H}_t \right]. \quad (19)$$

Define the Kalman gain matrix:

$$K_{t|t-1}^{(\mathcal{H}_t)} = P_{t|t-1}^{(\mathcal{H}_t)} Z_{s_t}' \left[F_{t|t-1}^{(\mathcal{H}_t)} \right]^{-1}. \quad (20)$$

Algorithm 3. *State Vector Smoothing for IMM and GPBN*

Step 0. Initialize the smoother by setting $r_{n|n-1}^{(\mathcal{H}_n)} = Z_{s_n}' \left[F_{n|n-1}^{(\mathcal{H}_n)} \right]^{-1} v_{n|n-1}^{(\mathcal{H}_n)}$, $N_{n|n-1}^{(\mathcal{H}_n)} = Z_{s_n}' \left[F_{n|n-1}^{(\mathcal{H}_n)} \right]^{-1} Z_{s_n}$, $\mathcal{H}_n \in \mathbb{H}_{N,n}$. Initialize $\alpha_{n|n}^{(\mathcal{H}_n)}$ and $P_{n|n}^{(\mathcal{H}_n)}$ by outputs of the corresponding filter at $t = n$.

Step 1. Compute the following auxiliary quantities for each regime, $i = 1, \ldots, h$, using recursion:

$$L_{t+1,t}^{(\mathcal{H}_t)} = T_{s_{t+1}} \left(I - K_{t|t-1}^{(\mathcal{H}_t)} Z_{s_t} \right)$$

$$r_{t|t-1}^{(\mathcal{H}_t)} = Z_{s_t}' \left[F_{t|t-1}^{(\mathcal{H}_t)} \right]^{-1} v_{t|t-1}^{(\mathcal{H}_t)} + \sum_{s_{t+1}=1}^{h} Q(s_t, s_{t+1}) L_{t+1,t}^{(\mathcal{H}_t)'} r_{t+1|t}^{(\mathcal{H}_{t+1})},$$

$$N_{t|t-1}^{(\mathcal{H}_t)} = Z_{s_t}' \left[F_{t|t-1}^{(\mathcal{H}_t)} \right]^{-1} Z_{s_t} + \sum_{s_{t+1}=1}^{h} Q(s_t, s_{t+1}) L_{t+1,t}^{(\mathcal{H}_t)'} N_{t+1|t}^{(\mathcal{H}_{t+1})} L_{t,t+1}^{(\mathcal{H}_t)}$$

for $t = n-1, n-2, \ldots, 1$.

Step 2. Compute the smoothed estimates of the state vector and the MSE matrix

$$\alpha_{t|n}^{(\mathcal{H}_t)} = \alpha_{t|t-1}^{(\mathcal{H}_t)} + P_{t|t-1}^{(\mathcal{H}_t)} r_{t|t-1}^{(\mathcal{H}_t)}, \quad P_{t|n}^{(\mathcal{H}_t)} = P_{t|t-1}^{(\mathcal{H}_t)} - P_{t|t-1}^{(\mathcal{H}_t)} N_{t|t-1}^{(\mathcal{H}_t)} P_{t|t-1}^{(\mathcal{H}_t)},$$

for $t = 1, \ldots, n-1$.

Use the smoothed probabilities, $\left\{ \mu_{t|n}^{(\mathcal{H}_t)} \right\}_{t=1:n-1}$, to compute the smoothed state vectors and MSE matrices:

$$\alpha_{t|n} = \sum_{\mathcal{H}_t} \mu_{t|n}^{(\mathcal{H}_t)} \alpha_{t|n}^{(\mathcal{H}_t)}, \quad P_{t|n} = \sum_{\mathcal{H}_t} \mu_{t|n}^{(\mathcal{H}_t)} P_{t|n}^{(\mathcal{H}_t)}.$$

See Appendix D for the detailed derivations.

The pseudocode below summarizes smoothing for regime probabilities and state means. The corresponding covariance recursions are omitted, as these are rarely computed in practice.

Pseudocode.

Implementation of smoothing for GPB*1*/IMM and GPB*2*

Index sets: time $t = 1,\ldots,n$; regimes $i, j, m, u \in \{1,\ldots,h\}$.

1. Inputs. Filter outputs $\{\alpha_{t|t-1}, P_{t|t-1}, v_{t|t-1}, F_{t|t-1}, K_t\}$, transition probabilities *Q*, updated regime probabilities $\mu_{t|t}$.

2. Outputs. Smoothed states $\alpha_{t|n}$ and smoothed regime probabilities $\mu_{t|n}$.

3. Initialization at $t = n$.

$r^{j}_{n|n-1} = Z'_{j,n} F^{-1}_{j,n} v^{j}_{n|n-1}$ (GPB1/IMM),

$r^{ij}_{n|n-1} = Z'_{j,n} F^{-1}_{ij,n} v^{ij}_{n|n-1}$ (GPB2).

4. Backward recursion for $t = n-1,\ldots,1$:

Smoothed probabilities (Kim, 1994):

$$\mu^{j}_{t|h} = \sum_{m=1}^{h} \mu^{m}_{t+1|h} \frac{\mu^{j}_{t|t} Q^{jm}}{\sum_{u=1}^{h} Q^{um} \mu^{u}_{t|t}}.$$

Smoothing transition:

$L^{jm}_{t+1,t} = T_{m,t+1}(I - K_{j,t} Z_{j,t})$ (GPB1/IMM),

$L^{ijm}_{t+1,t} = T_{m,t+1}(I - K_{i,j,t} Z_{j,t})$ (GPB2).

Adjoint update:

$r^{j}_{t|t-1} = Z'_{j,t} F^{-1}_{j,t} v^{j}_{t|t-1} + \sum_{m=1}^{h} Q^{jm} L^{jm'}_{t+1,t} r^{m}_{t+1|t}$ (GPB1/IMM),

$r^{ij}_{t|t-1} = Z'_{j,t} F^{-1}_{i,j,t} v^{ij}_{t|t-1} + \sum_{m=1}^{h} Q^{jm} L^{ijm'}_{t+1,t} r^{jm}_{t+1|t}$ (GPB2).

Smoothed states:

$$\alpha_{t|n}^{j} = \alpha_{t|t-1}^{j} + P_{t|t-1}^{j} r_{t|t-1}^{j} \quad \text{(GPB1/IMM)},$$

$$\alpha_{t|n}^{ij} = \alpha_{t|t-1}^{ij} + P_{t|t-1}^{ij} r_{t|t-1}^{ij} \quad \text{(GPB2)}.$$

5. Merging step.

$$\alpha_{t|n}^{j} = \sum_{i=1}^{h} \mu_{t-1|t}^{ij} \alpha_{t|n}^{ij}, \quad \text{(GPB2only)}$$

$$\alpha_{t|n} = \sum_{j=1}^{h} \mu_{t|n}^{j} \alpha_{t|n}^{j}. \quad \text{(all cases)}$$

Intuitively, our smoother follows the same principle as the disturbance–smoother for single–regime state–space models. The forward filter has already computed all the necessary one–step–ahead prediction errors, their variances $F_{t|t-1}$, and the associated Kalman gains $K_{t|t-1}$. Rather than recomputing large joint covariances of the entire state vector, the smoother reuses these filter outputs in a backward recursion. This has two main advantages. First, the computational cost is low: the additional backward pass involves only matrix–vector products of the same dimension as the state, with no need to form or invert large block matrices. Second, the procedure is numerically stable even when the measurement–error covariance is singular, a case where the Kim (1994) smoother for latent variables fails. This intuition underlies the recursions presented below and is what makes our smoother well suited to macroeconomic models with zero measurement errors.

4 Validating the Filters and Smoothers

We validate the performance of the filters and the new smoother using three complementary metrics: likelihoods, forecast errors (RMSEs), and hit rates.

Definitions. For each method *X* relative to the baseline method *B* we report:

$$\begin{aligned}
&\text{Log-likelihood difference}: && \Delta\ell = \ell^{X} - \ell^{B},\\
&\text{Relative RMSE(\%)}: && 100 \times \frac{\mathcal{R}^{X} - \mathcal{R}^{B}}{\mathcal{R}^{B}},\\
&\text{Hitrate difference(pp)}: && 100 \times (\text{Hit}^{X} - \text{Hit}^{B}).
\end{aligned}$$

We use GPB*2* as the baseline. Here $\mathcal{R}^{X}$ denotes the RMSE under the method *X*, defined as follows.

For a latent variable α_t observed on a sample of length *n*, the RMSE for method *X* is

$$\mathcal{R}_{t|s}^{X} = \frac{1}{n_{sim}} \sum_{i=1}^{n_{sim}} \sqrt{\frac{1}{n} \sum_{t=1}^{n} \left(\frac{\alpha_t - \alpha_{t|s}^{X}}{\alpha_{ss}} \right)^2},$$

where $s = t$ denotes the update error and $s = n$ the smoothing error. The normalization by the steady-state level α_{ss} is applied to state variables, but not to regime probabilities.

The proportional gain from smoothing is then defined as

$$\text{Gain} = 1 - \frac{\mathcal{R}_{t|n}}{\mathcal{R}_{t|t}}.$$

Likelihoods provide a rigorous benchmark because each filter delivers an approximation to the full mixture likelihood. The mixture likelihood integrates over all regime histories and is therefore the theoretical upper bound. Approximating this mixture inevitably yields lower log-likelihood values, but these are directly comparable across filters. Hence, differences in $\ell_X - \ell_B$ measure the relative quality of the approximation, making log-likelihoods the natural criterion for ranking filters.

RMSEs complement this by focusing on how accurately the filter or smoother recovers the latent state variables that drive the model. Because the artificial data are generated from a known DGP, we can directly compare the recovered paths with the true latent states. Whereas likelihoods assess the internal statistical efficiency of the approximation, RMSEs measure the external accuracy of state recovery, the feature most relevant for forecasting and policy analysis.

Hit rates measure how accurately the filter classifies the latent regime at each date, by comparing the regime with the highest posterior probability to the true regime. This provides a simple and widely reported metric for regime identification. We therefore compute and report hit rates alongside likelihoods and RMSEs. Appendix F reports complete results; here we focus on the most informative comparisons.

4.1 Model and Simulation Design

Model.

We use the medium-scale New Keynesian DSGE model of Fernandez-Villaverde, Guerron-Quintana, and Rubio-Ramirez (2015) (hereafter FGR2015), modified to incorporate two independent Markov-switching processes. Households derive utility from habit-adjusted consumption and leisure and supply differentiated labor under Calvo wage frictions. Firms produce output using capital, labor, and neutral technology with Calvo price setting. Capital accumulation features embodied investment-specific technology. The monetary policy follows a Taylor-type rule:

$$\frac{r_t}{r_{ss}} = \left(\frac{r_{t-1}}{r_{ss}}\right)^{\gamma_r(S_{P,t})} \left[\left(\frac{\pi_t}{\pi_{\text{targ}}}\right)^{\gamma_\pi(S_{P,t})} \left(\frac{Y_{d,t}}{\lambda_{yd} Y_{d,t-1}}\right)^{\gamma_y(S_{P,t})}\right]^{1-\gamma_r(S_{P,t})} \exp\left(\sigma_\xi(S_{V,t})\varepsilon_{\xi,t}\right), \quad (21)$$

where γ-parameters are high (hawkish) in policy state $S_{P,t} = 1$ and low (dovish) in $S_{P,t} = 2$. Another process, $S_{V,t}$, governs shock volatilities: in the low-volatility state, the variances match the FGR2015 estimates, while in the high-volatility state they are doubled (for all structural shocks, including policy). Full model details are in Appendix E.

Parameterization.

The structural parameters follow FGR2015 (see column (1) in Table E1, Appendix); in addition, we introduce regime switching in both policy and volatility. For the benchmark, transition matrices are assumed *symmetrical*, with identical persistence ρ on the main diagonal.

Simulation design.

The model is solved and simulated using the functional iteration algorithm in the `RISE` toolbox. We generate 500 artificial samples of length 1,000. Our observables include output growth, price inflation, wage inflation, the Federal Funds rate, and the relative price of investment; latent variables include consumption, capital, output, and the regime states. To assess filter–smoother performance, we evaluate (i) the first 300 observations (a "short sample"), (ii) the last 500 observations (where initial conditions fade), and (iii) the full 1,000 observations.

4.2 Ranking of Filters and Smoothers

Illustration: what smoothing does.

Before turning to formal rankings, it is useful to illustrate how smoothing modifies the signals extracted by the filter. Figure 1 shows the output of a representative filter in both updating and smoothing modes, for a sample of 300 observations.

The top two panels display regime probabilities for volatility and policy. Here regime probabilities are smoothed with the backward recursion of Kim (1994) (generalized in Appendix C), which is widely used in the literature. As expected, the smoothed probabilities are less noisy than their filtered counterparts, but initial-condition effects are minimal and the qualitative picture is unchanged (see Appendix F.2, Figure F.1 for further confirmation). This illustrates why probabilities are not the main focus of our analysis: the benefits of smoothing in this dimension are well understood.

The bottom two panels display economic variables: consumption, a fast-moving forward-looking household decision variable, and capital, a slow-moving predetermined state variable. Here we apply our new smoother. In contrast to the probabilities, large initial-condition effects are visible in the filtered estimates, especially for capital. The smoother uses future as well as past

information to correct these distortions, producing trajectories that track the data more closely and are visibly more accurate.

This simple illustration motivates the validation analysis that follows. For probabilities, smoothing delivers expected refinements without altering the overall picture. For economic variables, by contrast, smoothing can substantially improve accuracy when initialization errors are present. It is this dimension—the recovery of latent states and observables—that underpins our formal comparisons in the next subsections.

Likelihoods.

Table 1 reports log-likelihood differences ($\Delta\ell$) relative to GPB*2*. IMM and GPB*3* are statistically indistinguishable from GPB*2* across all samples, while GPB*1* is consistently and significantly worse. This ranking is stable: GPB*1* is dominated, while the other three filters cannot be separated on likelihood grounds. Hence GPB*2* serves as the benchmark in what follows.

Forecast errors (RMSEs).

Table 2 reports relative RMSEs for consumption (decision variable, flow) and capital (predetermined state, stock). GPB*1* again performs consistently worse, while IMM and GPB*3* are not distinguishable from GPB*2* in the updates. For the smoother, the ranking depends on sample length. With 1000 observations IMM already achieves significantly lower RMSEs for consumption and output, while capital is similar to GPB*2*. In the last 500 observations, where initial-condition effects fade, IMM smoothing also improves capital significantly. Full results for all variables (Table F.1 in Appendix F.2) confirm that update-stage RMSEs for IMM and GPB*2* are statistically indistinguishable. This makes it clear that subsequent performance differences arise from the smoothing step rather than from the underlying filter updates.

Accuracy gains from smoothing.

Table 3 reports proportional reductions in RMSE when moving from filtering to smoothing, defined as $1-(\mathcal{R}_{t|n}/\mathcal{R}_{t|t})$. Gains are substantial, ranging from around 15% for capital to about 25% for consumption in medium-length samples. Table F.2 in Appendix F.2 reports the gains for the full set of variables, for which the improvements are even larger for some observables. The magnitude of the gain is broadly stable across sample lengths and smaller in the last 500 observations, where initialization errors are absent.

Comparing filters, IMM and GPB*2* achieve similar improvements, with IMM showing significantly larger gains for output and consumption once initialization effects fade. Capital is usually indistinguishable, with a modest IMM advantage in the last 500 observations. Thus, the main effect comes from the smoothing step itself, with IMM having a systematic edge in longer samples.

Hit rates.

Tables F.3–F.4 in Appendix F.3 reports regime classification hit rate differences $100\times(\text{Hit}^{X}-\text{Hit}^{B})$ based on updated and smoothed probabilities. Differences across filters are numerically small and add little beyond what is already captured by likelihoods and RMSEs.

Computational burden.

Practical performance also depends on computational cost. IMM runs faster than GPB2 because it propagates *h* histories rather than h^2. Table 4 reports indicative run times for the two filters, both alone and with their smoothers. [8] Although the absolute times appear small, repeated evaluation during estimation quickly accumulates. For example, a model with 300 quarterly observations and around 30 parameters may require hours per optimization run, and subsequent MCMC simulations can take weeks. Relative comparisons are therefore more informative: IMM is roughly three times faster than GPB*2*, while higher-order GPB filters escalate sharply in cost. Balancing accuracy and speed, IMM emerges as the most practical choice. This advantage is relevant both in maximum likelihood estimation, where the filter is called at every parameter trial, and even more in Bayesian estimation contexts, where it must be evaluated thousands of times in posterior sampling. In contrast, the novelty of our smoother lies not in speed but in robustness: it reuses filter outputs in a single backward pass and remains stable even in the zero–measurement–error case where existing smoothers fail or are infeasible.

Summary.

Likelihood comparisons show that GPB*2*, IMM, and GPB*3* are statistically equivalent, while GPB*1* is clearly dominated. As discussed in Appendix F.1, this reflects the fact that all these filters approximate the full mixture likelihood by collapsing regime histories through moment matching, a design not intended to maximize likelihood. In addition, in a linear–Gaussian model, retaining longer histories (as in GPB*3*) adds little once collapsing has occurred, so GPB*2* and GPB*3* deliver virtually identical likelihoods in all samples.

RMSE results confirm this ranking and demonstrate the role of smoothing. Smoothing consistently improves accuracy by 15–35% across variables and samples, with larger gains for fast-moving decision variables than for slow-moving state variables. Importantly, the proportional gains are not uniform across filters: they are largest for IMM and GPB*1*, which collapse regime histories earlier in the forward pass. This is intuitive: aggressive mixing in the forward pass leaves more scope for the backward recursion to "unmix" errors using future information, so the smoother eliminates a larger share of the update RMSE. By contrast, higher-order GPB filters start closer to the truth and therefore gain less from smoothing. Appendix F.2 develops this mechanism in detail, showing how variance inflation, regime misclassification, and initial-condition sensitivity combine to give low-order filters more scope for improvement.

From a practical perspective, IMM combines these accuracy gains with substantially lower computational cost: about three times faster than GPB*2*, while higher-order GPB filters become

prohibitively expensive. Taken together, these results validate the filters, quantify the magnitude of smoothing gains, and explain why the IMM filter–smoother pair emerges as the most practical choice for applied work.

4.3 Robustness to Alternative Scenarios

All robustness checks are structured around two dimensions. The first is always the persistence of the latent regime processes, which we vary across $\rho \in \{0.80, 0.90, 0.95\}$. The second dimension differs by exercise. In Section 4.2 we paired regime persistence with sample length to study how the rankings evolve with longer histories. Here we fix the sample length at 300 observations and instead pair regime persistence with (i) a comparison of filter rankings (IMM vs. GPB*2*), (ii) alternative persistence of structural shocks and (iii) misspecification of the regime process. This design ensures that our findings are not driven by near–absorbing regimes or by a particular feature of the baseline calibration. Throughout, $\rho = 0.90$ serves as the default case, representing moderately persistent but non-absorbing regimes.

Regime persistence and filter rankings.

We first ask whether the relative performance of IMM and GPB*2* depends on the persistence of regimes. Table 5 compares the two smoothers at $\rho \in \{0.80, 0.90, 0.95\}$, reporting both relative RMSE differences (IMM–GPB*2*, negative = IMM more accurate) and proportional RMSE reductions from smoothing ("gains"). Likelihood comparisons are unchanged across ρ—IMM and GPB*2* remain statistically indistinguishable on $\Delta\ell$—so we omit those rows for brevity (see Appendix F.4, Table F.5).

For consumption, IMM is more accurate at high persistence ($\rho = 0.95$), while GPB*2* dominates at low persistence ($\rho = 0.80$); at $\rho = 0.90$ the methods are broadly indistinguishable. Capital shows the opposite pattern, with GPB2 favored except at the highest persistence. Thus, smoothing yields large and robust accuracy improvements at all persistence levels, so the benefit does not rely on near-absorbing states. The choice between IMM and GPB*2* is persistence-dependent: IMM is preferred in high-persistence regimes, while GPB*2* is preferable in low-persistence regimes.

Shock persistence and smoothing performance.

We next ask whether the smoother's performance changes when structural shocks are more persistent. Specifically, we increase the autoregressive coefficients of the preference and labor-supply shocks by halving their distance from a unit root: $\rho^{\text{new}} = 1 - \frac{1}{2}(1 - \rho^{\text{old}})$. This raises the persistence of the preference shock from 0.118 to 0.559 and the labor-supply shock from 0.933 to 0.967. Because the effect is large and transparent in levels, we focus on the IMM smoother only and report RMSE levels alongside proportional smoothing gains in Table 6.

Two conclusions follow.

First, fast-moving variables such as consumption remain robust: RMSEs stay low (around 0.08) and smoothing gains large (around 25%), essentially unchanged from the benchmark.

Second, the slow-moving stock variable, capital, is more sensitive: RMSEs rise five- to sixfold (to 8–9) and smoothing gains shrink markedly, especially at $\rho = 0.95$ where they nearly vanish (0.081).

Overall, higher shock persistence primarily undermines the recovery of slow-moving variables while leaving fast-moving variables relatively unaffected. This reinforces our earlier message: smoothing is valuable across calibrations, but its relative effectiveness depends on both variable dynamics and shock persistence.

Misspecification of policy regimes.

We next consider the effect of misspecifying the number of policy regimes. Table 7 compares RMSE levels and smoothing gains for the correctly specified IMM smoother (two policy regimes) with a misspecified model that assumes only one.

Two patterns emerge.

First, fast-moving variables such as consumption remain robust: RMSE levels differ little between specifications (0.035–0.042) and smoothing gains remain stable at 25–35% across persistence values.

Second, the slow-moving state variable, capital, is much more vulnerable. Under the correct model RMSEs remain stable around 1.4–1.6, but under misspecification they rise substantially—to nearly 2.8 at $\rho = 0.95$. The smoother still improves accuracy within the misspecified model, but the proportional gain is smaller (0.107 vs. 0.166 under the true specification).

Figure 2 illustrates these results when policy persistence is high ($\rho = 0.95$). Consumption closely tracks the data under both specifications, while capital drifts away under misspecification. The smoother corrects some of this drift, but not fully.

Overall, the smoother delivers meaningful accuracy gains even when policy regimes are misspecified, but the benefits are strongest when the regime structure is correctly specified.

Misspecification of volatility regimes.

A complementary exercise rules out volatility switching altogether, assuming shocks are always low- or always high-volatility. This misspecification delivers artificially low RMSEs and inflated smoothing gains, since regime transitions are mechanically excluded. For example, output RMSEs fall from about 0.040 under the true model to 0.023–0.025 under misspecification, while proportional smoothing gains rise from about 0.33 to more than 0.60.

These distortions appear under low- and high-volatility misspecifications and are robust across all persistence values ($\rho = 0.80, 0.90, 0.95$). The bias is structural: ignoring volatility regimes

makes the smoother appear much more effective than it truly is. Full results are reported in Table F.9 in Appendix F.8.

Particle–filter benchmark.

As an additional robustness check, Appendix G benchmarks the IMM and GPB*2* filters, together with our new smoother, against the Auxiliary Particle Filter (APF) of Pitt and Shephard (1999) and Doucet et al. (2000), extended with a Rao–Blackwellized smoother. To make this comparison feasible, we use a smaller RBC–style model rather than the full New Keynesian system, while still retaining both predetermined (capital) and forward–looking (consumption) variables. The results show that our smoother brings IMM and GPB*2* accuracy into line with the particle benchmark in terms of RMSEs, while the filter likelihoods confirm that IMM performs on par with GPB*2* and with the APF, extending the results in Kim and Kang (2019). At the same time, the APF requires extremely large particle sets before converging, which makes Gaussian–mixture filters combined with our smoother a more practical tool for macroeconomic applications.

5 Empirical Application

We now investigate the practicality of the IMM filter with its smoother. [9] We estimate a modified version of the FGR2015 model using the same data (1959Q2–2013Q4; see Table E.4 in Appendix E). Priors are relatively wide, and we use the *Artificial Bee Colony* algorithm (Karaboga and Basturk, 2007) for global optimization.

Table 8 reports estimated transition probabilities and policy parameters; the remaining parameters are in Table E.3 (Appendix E). Both regime-switching processes are highly persistent, so the model assigns high probability to staying in the same regime. In particular, the policy process implies only a 2% probability of leaving the hawkish state.

Panels A and B of Figure 3 show smoothed probabilities of the dovish policy state and the high volatility state, computed according to Kim (1994). We used the IMM at the estimation stage and six different filters at the filtering stage. The lines from these six filters are nearly indistinguishable: GPB filters of order from 2 to 5 give almost identical results and are extremely close to the IMM. Thus, higher-order GPB filters add little compared to the GPB*2* filter, and IMM and GPB*2* deliver practically identical regime probabilities. GPB*1* is less accurate but still captures the main events.

Panel A shows the probability of being in the dovish policy state. The inferred states align with major postwar macroeconomic episodes: the Great Inflation, Volcker Disinflation, Great Moderation, Global Financial Crisis, and Zero Lower Bound (ZLB) of 2008 episodes. We did not impose a ZLB regime; instead, the model identifies this period as dovish.

Panel B shows the high-volatility state, with probabilities elevated around most recessions. The pre-1990s period appears to have larger shocks than the recent past.

Panels C and D extend the dataset to 1947Q2–2023Q3 (see Appendix E). The longer span improves identification of the Great Inflation and includes the post-Covid period. Here we report results only for the IMM filter–smoother pair. The patterns are consistent with Panels A and B: the economy is in the dovish state during the ZLB of 2008, shifts hawkish about a year after lift-off, and remains dovish through the Covid-19 pandemic until 2023Q1, after which strong anti-inflation measures appear. The post-Covid period is also characterized by relatively large shocks.

6 Conclusions

This paper presents improved multiple-regime filtering techniques and a robust latent-variable smoother for regime-switching state-space models. The novel recursion-based smoother is designed for applications with zero measurement error—common in DSGE analysis—is computationally stable, and can be paired with a wide range of filters.

To assess its performance, we compared IMM and GPB filters in simulations with a prototypical New Keynesian DSGE model. The IMM filter is about three times faster and at least as accurate as the GPB*2* (Kim–Nelson) filter, and combining it with our new smoother improves recovery of latent variables by approximately 25%. An empirical application with long U.S. time series further confirms the practical relevance of these tools.

To make them accessible in applied research, we provide a comprehensive toolkit implementing this methodology. The code, which accommodates extensions such as time-varying states and parameters, endogenous regime switches, and missing data, is available in the RISE© toolbox for MATLAB©.

Notes

[1] See, *inter alia*, Davig and Doh, 2014, Chang, Maih, and Tan, 2021, Chen, Leeper, and Leith, 2022.

[2] RISE stands for "Rationality in Switching Environments". The toolbox and documentation are available at https://github.com/jmaih/RISE_toolbox

[3] See Chapter 13 of Frühwirth-Schnatter (2006) for a comprehensive review of approximate filters for switching state space models.

[4] In reference to dimension reduction, we use *mixing* in the description of the IMM and *collapsing* in the description of the GPB following the convention in the literature.

[5] See Table A.1 in Appendix A.

[6] See also Table A.1 in the Appendix A.

[7] See e.g. Helmick et al. (1995), Weinert (2001) and Balenzuela et al. (2022).

[8] Times should be interpreted as indicative only, since RISE implementations include overhead for general features such as time-varying states, missing observations, and occasionally-binding constraints; see Appendix H.

[9] Some earlier studies employed the IMM filter without smoothing. For example, Binning and Maih (2015) adapted non-linear filters to the multiple-regime setting, Bjørnland, Larsen, and Maih (2018) studied oil price shocks and macroeconomic instability, and Leith, Kirsanova, Machado, and Ribeiro (2025) analyzed U.S. monetary and fiscal policy. Beyond these, we are aware only of Liu, Wang, and Zha (2013), which apparently used the IMM filter to study land-price dynamics.

References

Balenzuela, M. P., A. G. Wills, C. Renton, and B. Ninness (2022). A new smoothing algorithm for jump Markov linear systems. *Automatica 140*, 110218.

Binning, A. and J. Maih (2015). Sigma point filters for dynamic nonlinear regime switching models. Working Paper 2015/10, Norges Bank.

Bjørnland, H. C., V. H. Larsen, and J. Maih (2018). Oil and Macroeconomic (In)stability. *American Economic Journal: Macroeconomics 10*(4), 128–51.

Blom, H. A. (1984). An efficient filter for abruptly changing systems. In *The 23rd IEEE Conference on Decision and Control*, pp. 656–658.

Blom, H. A. and Y. Bar-Shalom (1988). The interacting multiple model algorithm for systems with Markovian switching coefficients. *IEEE transactions on Automatic Control 33*(8), 780–783.

Chang, Y., J. Maih, and F. Tan (2021). Origins of Monetary Policy Shifts: A New Approach to Regime Switching in DSGE Models. *Journal of Economic Dynamics and Control 133*, 104235.

Chen, X., E. M. Leeper, and C. Leith (2022). Strategic Interactions in U.S. Monetary and Fiscal Policies. *Quantitative Economics 13*(2), 593–628.

Davig, T. and T. Doh (2014). Monetary Policy Regime Shifts and Inflation Persistence. *The Review of Economics and Statistics 96*(5), 862–875.

De Jong, P. (1988). A Cross-Validation Filter for Time Series Models. *Biometrika 75*(3), 594–600.

Doucet, A., S. Godsill, and C. Andrieu (2000). On sequential monte carlo sampling methods for bayesian filtering. *Statistics and Computing 10*(3), 197–208.

Durbin, J. and S. J. Koopman (2012). *Time Series Analysis by State Space Methods* (2 ed.), Volume 38 of *Oxford Statistical Science Series*. Oxford: Oxford University Press.

Fernandez-Villaverde, J., P. A. Guerron-Quintana, and J. Rubio-Ramirez (2015). Estimating Dynamic Equilibrium Models with Stochastic Volatility. *Journal of Econometrics 185*(1), 216–229.

Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*. New York: Springer.

Helmick, R. E., W. D. Blair, and S. A. Hoffman (1995). Fixed-Interval Smoothing for Markovian Switching Svstems. *IEEE Transactions on Information Theory 41*(6), 1845–1855.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of basic Engineering 82*(1), 35–45.

Karaboga, D. and B. Basturk (2007). A powerful and efficient algorithm for numerical function optimization: artificial bee colony (ABC) algorithm. *Journal of Global Optimization 39*, 459–471.

Kim, C. and C. Nelson (1999). *State-space Models with Regime Switching: Classical and Gibbs-sampling Approaches with Applications*. MIT Press.

Kim, C.-J. (1994). Dynamic linear models with Markov-switching. *Journal of Econometrics 60*(1-2), 1–22.

Kim, Y. M. and K. H. Kang (2019). Likelihood inference for dynamic linear models with markov switching parameters: on the efficiency of the kim filter. *Econometric Reviews 38*(10), 1109–1130.

Leith, C., T. Kirsanova, C. Machado, and A. P. Ribeiro (2025). (Re)Evaluating recent macroeconomic policy in the US. *European Economic Review 178*, 105091.

Liu, Z., P. Wang, and T. Zha (2013, May). Land Price Dynamics and Macroeconomic Fluctuations. *Econometrica 81*(3), 1147–1184.

Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association 94*(446), 590–599.

Weinert, H. L. (2001). *Fixed Interval Smoothing for State Space Models*. New York: Springer.

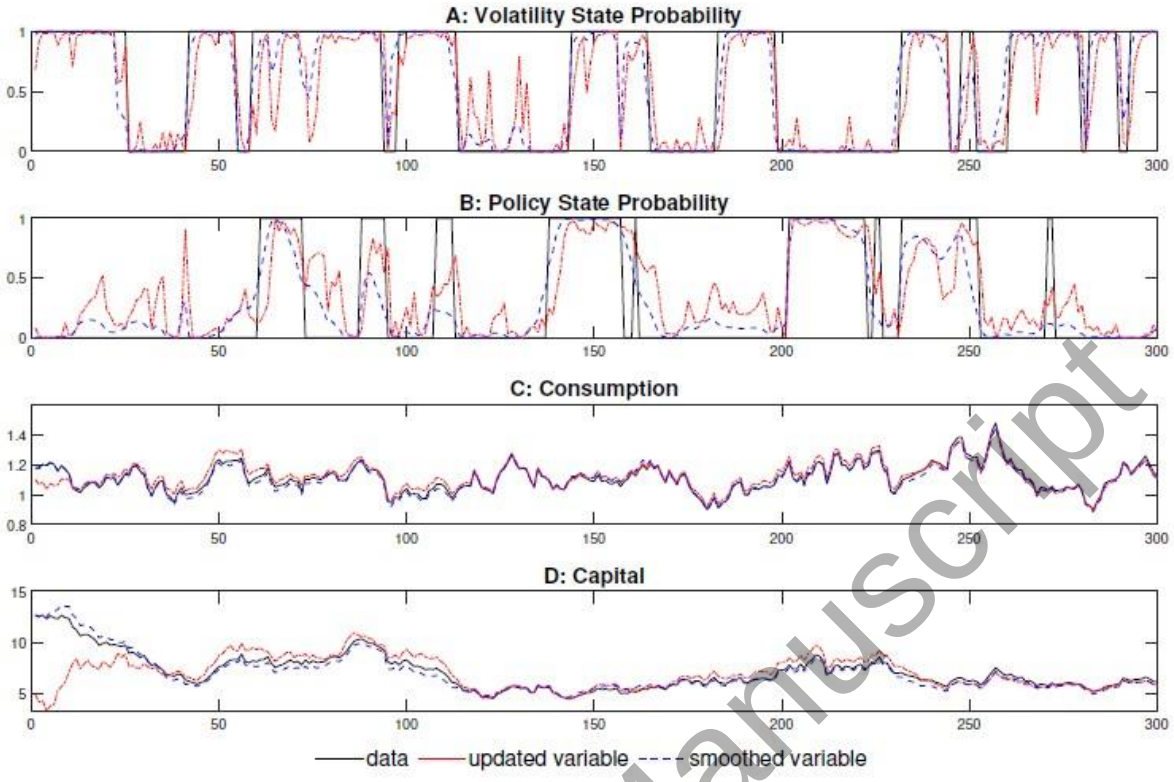

Figure 1: Updating and smoothing. Panels A–B: regime probabilities for volatility and policy, using Kim's probability smoother. Panels C–D: economic variables (consumption and capital) using our new smoother.

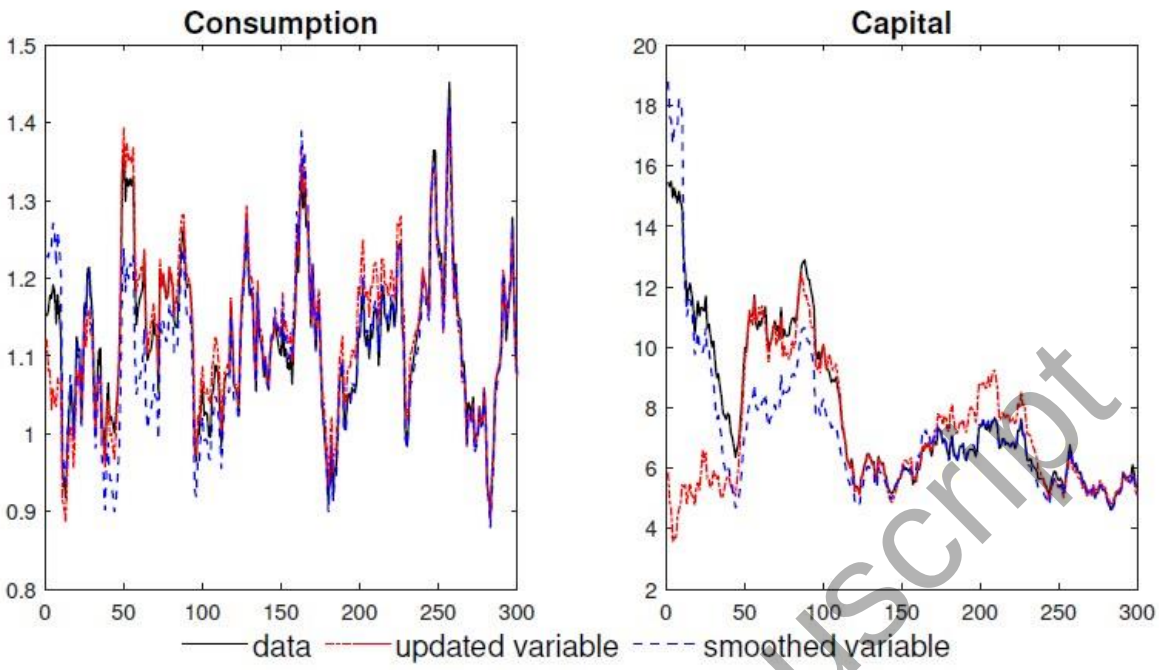

Figure 2: Illustration of misspecifying the number of policy regimes.

Note: Results shown for $\rho = 0.95$ with a 300-observation sample. The correctly specified model includes two policy regimes; the misspecified model assumes only one.

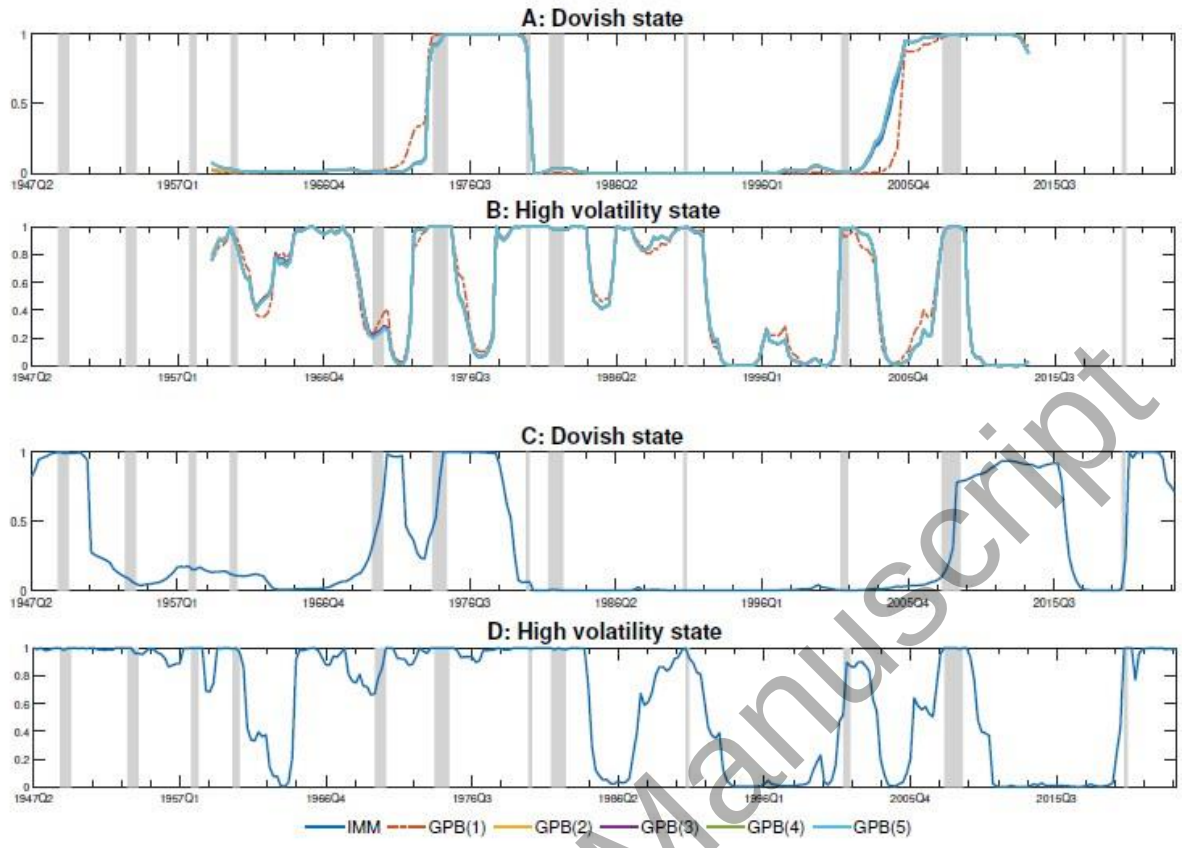

Figure 3: Smoothed State Probabilities.

Table 1: Log-likelihood differences relative to GPB*2*.

Sample	IMM	GPB*1*	GPB*3*
50 obs	0.00092 (1.10)	–0.114 (-5.75)	–0.00014 (-0.51)
100 obs	0.00085 (0.85)	–0.211 (-7.17)	–0.00011 (-0.28)
200 obs	0.00034 (0.25)	–0.406 (-9.69)	–0.00014 (-0.33)
300 obs	0.00021 (0.13)	–0.549 (-10.80)	–0.00013 (-0.30)
1000 obs	–0.00001 (-0.00)	–1.792 (-20.45)	–0.00019 (-0.44)
Last 500	–0.00057 (-0.35)	–0.857 (-15.11)	–0.00001 (-0.39)

Note: Parentheses show t-statistics. Bold t-statistics indicate significance at conventional levels. Positive values of $\Delta\ell$ imply method *X* outperforms GPB*2*.

Table 2: Relative RMSEs vs GPB*2* for key variables.

	Updates (1000 obs)			IMM smoother	
Variable	IMM	GPB*1*	GPB*3*	1000 obs	Last 500
Consumption	0.0020 (0.17)	1.8984 (5.60)	0.0030 (0.48)	–0.6146 (-5.35)	–0.6432 (-5.78)
Capital	0.0044 (0.21)	2.0606 (4.48)	0.0165 (0.94)	–0.0993 (-0.37)	–0.9086 (-5.44)

Note: Relative RMSEs are computed as $100\times(\mathcal{R}^{X}-\mathcal{R}^{GPB2})/\mathcal{R}^{GPB2}$. Parentheses report *t*-statistics; bold values indicate significance at conventional levels. Negative entries imply method *X* outperforms GPB*2*.

Table 3: RMSE improvements from smoothing (IMM vs GPB*2*).

	50 obs	100 obs	200 obs	300 obs	1000 obs	Last 500
Consumption						
IMM	0.226 (2.67)	0.226 (2.73)	0.248 (0.20)	0.241 (−1.36)	0.222 (−5.06)	0.064 (−3.75)
GPB*2*	0.229	0.228	0.248	0.240	0.217	0.058
Capital						
IMM	0.079 (0.98)	0.121 (−0.04)	0.169 (−0.19)	0.166 (−0.71)	0.156 (−0.83)	0.070 (−2.04)
GPB*2*	0.090	0.125	0.173	0.168	0.155	0.062

Note: Entries are proportional reductions in RMSE from applying the smoother. IMM values include t-statistics (in parentheses) for the difference relative to GPB*2*. Significant negative t-statistics indicate IMM delivers larger accuracy gains than GPB*2*. For full table see Table F.2

Table 4: Indicative computational cost ($n = 1000$ observations).

	Computation time (sec)		Relative speed
	Updating only	Updating + smoothing	(GPB*2*=1)
IMM	0.27	1.49	0.28
GPB*2*	1.38	2.59	1.00
GPB*3*	–	–	4.21
GPB*4*	–	–	17.74
GPB*5*	–	–	79.97

Note: The computations were carried out on a Ryzen 3950X with 64GB RAM using MATLAB R2022b. Relative speed shows the ratio of updating times, normalized to GPB*2*.

Table 5: Robustness of smoothing performance to regime persistence (300 observations).

	$\rho = 0.80$	$\rho = 0.90$	$\rho = 0.95$
Consumption			
Rel. RMSE (IMM–GPB*2*)	0.735 (**20.29**)	−0.184 (-1.72)	−1.532 (-7.82)
IMM gain	0.245 (**9.39**)	0.241 (-1.36)	0.248 (**−6.49**)
GPB*2* gain	0.249	0.240	0.236
Capital			
Rel. RMSE (IMM–GPB*2*)	0.312 (**5.64**)	0.219 (0.78)	1.438 (**2.15**)
IMM gain	0.160 (**1.97**)	0.166 (-0.71)	0.166 (-0.09)
GPB*2* gain	0.162	0.168	0.179

Note: Relative RMSE rows report IMM minus GPB2; "Gain" rows report proportional RMSE reductions from smoothing.

Table 6: Effect of higher shock persistence (300 obs).

	$\rho = 0.80$	$\rho = 0.90$	$\rho = 0.95$
Consumption			
RMSE (benchmark)	0.035	0.036	0.035
RMSE (high persistence)	0.082	0.082	0.079
Gain (benchmark)	0.245	0.241	0.248
Gain (high persistence)	0.253	0.248	0.237
Capital			
RMSE (benchmark)	1.571	1.556	1.410
RMSE (high persistence)	8.355	8.702	8.447
Gain (benchmark)	0.160	0.166	0.166
Gain (high persistence)	0.168	0.134	0.081

Notes: Entries are RMSE levels and proportional smoothing gains. Benchmark refers to the baseline calibration in Section 4.4.2.

Table 7: Effect of misspecifying the number of policy regimes (300 obs).

	$\rho = 0.80$	$\rho = 0.90$	$\rho = 0.95$
Consumption			
True (IMM smoother)	0.035	0.036	0.035
Misspecified (1 regime)	0.035	0.037	0.042
Gain (true)	0.245	0.241	0.248
Gain (misspec.)	0.249	0.240	0.237
Capital			
True (IMM smoother)	1.571	1.556	1.410
Misspecified (1 regime)	1.589	1.686	2.797
Gain (true)	0.160	0.166	0.166
Gain (misspec.)	0.165	0.174	0.107

Notes: Entries report RMSE levels and proportional smoothing gains for the IMM smoother under the correct two-regime specification and a misspecified one-regime model.

Table 8: Estimation of parameters that govern the two Markov processes

Transition matrices	
Shocks	**Policy**
$P_s = \begin{bmatrix} 0.93905 & 0.060946 \\ 0.04625 & 0.95375 \end{bmatrix}$	$P_p = \begin{bmatrix} 0.98396 & 0.016039 \\ 0.043428 & 0.95657 \end{bmatrix}$
Taylor Rule Parameters:	
Hawkish Feedback γ_π	1.6574
Dovish Feedback γ_π	0.93984