

# ADDITIONAL RESULTS

## Inflation and Professional Forecast Dynamics: An Evaluation of Stickiness, Persistence, and Volatility\*

INTENDED FOR ONLINE PUBLICATION ONLY

Elmar Mertens<sup>†</sup>  
Deutsche Bundesbank

James M. Nason<sup>‡</sup>  
NC State University

April 26, 2020

### Abstract

This online appendix provides additional estimates and results of a Monte Carlo study to complement the results shown in our paper.

---

\*The views herein are those of the authors and do not represent the views of the Deutsche Bundesbank or the Eurosystem.

<sup>†</sup>*email:* elmar.mertens@bundesbank.de, *address:* Deutsche Bundesbank, Central Office-Research Centre, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main, Germany.

<sup>‡</sup>*email:* jmnason@ncsu.edu, *address:* Department of Economics, Campus Box 8110, NC State University, Raleigh, NC 27695 and Centre for Applied Macroeconomic Analysis. Jim Nason thanks the Jenkins Family Economics Fund at North Carolina State University for financial support.

# Contents

<b>R.1</b>	<b>Model-based inflation forecasts and the SPF</b>	<b>6</b>
R.1.1	SPF term structure of inflation forecasts . . . . .	6
R.1.2	Forecasting Inflation with Our State Space Models . . . . .	8
<b>R.2</b>	<b>Detailed state-space model results</b>	<b>13</b>
R.2.1	Results for the four model variants . . . . .	13
R.2.2	Results for case without noise component in inflation . . . . .	30
R.2.3	Log MDD over time and for different numbers of particles . . . . .	48
R.2.4	Relative effective sample sizes . . . . .	53
<b>R.3</b>	<b>An alternative specification for time variation in <math>\lambda_t</math></b>	<b>56</b>
<b>R.4</b>	<b>An alternative forecast error variance decomposition</b>	<b>64</b>
<b>R.5</b>	<b>SPF predictions and trend inflation uncertainty</b>	<b>66</b>
<b>R.6</b>	<b>Univariate estimates from the Stock-Watson UCSV model</b>	<b>68</b>
<b>R.7</b>	<b>Results with CPI data</b>	<b>73</b>
R.7.1	Results for the four model variants . . . . .	73
R.7.2	CPI results for case without noise component in inflation . . . . .	91
R.7.3	Log MDD for different numbers of particles (CPI) . . . . .	109
R.7.4	CPI data: relative effective sample sizes . . . . .	112
R.7.5	Forecasting CPI Inflation with Our State Space Models . . . . .	112
<b>R.8</b>	<b>Monte Carlo Study</b>	<b>117</b>
R.8.1	Simulated bias in estimates . . . . .	117
R.8.2	Model detection . . . . .	122

## List of Tables

R.1	Inflation Forecast Comparison . . . . .	11
R.2	Inflation Forecast Comparison (models with noise in inflation) . . . . .	12
R.3	Parameter Estimates and MDDs w/o noise in inflation . . . . .	31
R.4	Parameter Estimates and MDDs (CPI) . . . . .	74
R.5	Parameter Estimates and MDDs w/o noise in inflation (CPI) . . . . .	92
R.6	Inflation Forecast Comparison (CPI data) . . . . .	115
R.7	Inflation Forecast Comparison (CPI data, models with noise in inflation) . . . . .	116

## List of Figures

R.1	Realized Inflation and SPF Predictions . . . . .	7
R.2	Term structures of SPF data, SI and RE forecasts . . . . .	9
R.3	$\mathcal{M}_0$ : Trend and Gap Inflation . . . . .	14
R.4	$\mathcal{M}_0$ : Static Parameters . . . . .	15
R.5	$\mathcal{M}_0$ : Stochastic Volatility in Trend and Gap Inflation . . . . .	16
R.6	$\mathcal{M}_0$ : Time-Varying SI Parameter . . . . .	17
R.7	$\mathcal{M}_1$ : Trend and Gap Inflation . . . . .	18
R.8	$\mathcal{M}_1$ : Static Parameters . . . . .	19
R.9	$\mathcal{M}_1$ : Stochastic Volatility in Trend and Gap Inflation . . . . .	20
R.10	$\mathcal{M}_2$ : Trend and Gap Inflation . . . . .	21
R.11	$\mathcal{M}_2$ : Static Parameters . . . . .	22
R.12	$\mathcal{M}_2$ : Stochastic Volatility in Trend and Gap Inflation . . . . .	23
R.13	$\mathcal{M}_2$ : Time-Varying SI Parameter . . . . .	24
R.14	$\mathcal{M}_2$ : Time-Varying AR Coefficient in the Inflation Gap Process . . . . .	25
R.15	$\mathcal{M}_3$ : Trend and Gap Inflation . . . . .	26
R.16	$\mathcal{M}_3$ : Static Parameters . . . . .	27
R.17	$\mathcal{M}_3$ : Stochastic Volatility in Trend and Gap Inflation . . . . .	28
R.18	$\mathcal{M}_3$ : Time-Varying AR Coefficient in the Inflation Gap Process . . . . .	29
R.19	$\mathcal{M}_0$ (w/o noise): Trend and Gap Inflation . . . . .	32
R.20	$\mathcal{M}_0$ (w/o noise): Static Parameters . . . . .	33
R.21	$\mathcal{M}_0$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation . . . . .	34
R.22	$\mathcal{M}_0$ (w/o noise): Time-Varying SI Parameter . . . . .	35
R.23	$\mathcal{M}_1$ (w/o noise): Trend and Gap Inflation . . . . .	36
R.24	$\mathcal{M}_1$ (w/o noise): Static Parameters . . . . .	37
R.25	$\mathcal{M}_1$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation . . . . .	38
R.26	$\mathcal{M}_2$ (w/o noise): Trend and Gap Inflation . . . . .	39
R.27	$\mathcal{M}_2$ (w/o noise): Static Parameters . . . . .	40
R.28	$\mathcal{M}_2$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation . . . . .	41
R.29	$\mathcal{M}_2$ (w/o noise): Time-Varying SI Parameter . . . . .	42
R.30	$\mathcal{M}_2$ (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process . . . . .	43
R.31	$\mathcal{M}_3$ (w/o noise): Trend and Gap Inflation . . . . .	44
R.32	$\mathcal{M}_3$ (w/o noise): Static Parameters . . . . .	45
R.33	$\mathcal{M}_3$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation . . . . .	46

R.34	$\mathcal{M}_3$ (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process	47
R.35	log MDD over time	49
R.36	log MDD over time (w/o noise)	50
R.37	log MDD for different numbers of particles	51
R.38	log MDD for different numbers of particles (w/o noise)	52
R.39	Relative ESS	54
R.40	Relative ESS (w/o noise)	55
R.41	$\mathcal{M}_{2,\lambda_x}$ : Trend and Gap Inflation	59
R.42	$\mathcal{M}_{2,\lambda_x}$ : Static Parameters	60
R.43	$\mathcal{M}_{2,\lambda_x}$ : Stochastic Volatility in Trend and Gap Inflation	61
R.44	$\mathcal{M}_{2,\lambda_x}$ : Time-Varying SI Parameter	62
R.45	$\mathcal{M}_{2,\lambda_x}$ : Time-Varying AR Coefficient in the Inflation Gap Process	63
R.46	Alternative $R^2$ measure of inflation persistence	65
R.47	Uncertainty around trend inflation estimates	67
R.48	Stock-Watson UCSV model: Inflation Trend and Gap	69
R.49	Stock-Watson UCSV model: SV in Inflation Trend and Gap	70
R.50	Stock-Watson UCSV model: Parameter estimates	71
R.51	Stock-Watson UCSV model: Effective Sample Size	72
R.52	$\mathcal{M}_0$ : Trend and Gap Inflation	75
R.53	$\mathcal{M}_0$ : Static Parameters	76
R.54	$\mathcal{M}_0$ : Stochastic Volatility in Trend and Gap Inflation	77
R.55	$\mathcal{M}_0$ : Time-Varying SI Parameter	78
R.56	$\mathcal{M}_1$ : Trend and Gap Inflation	79
R.57	$\mathcal{M}_1$ : Static Parameters	80
R.58	$\mathcal{M}_1$ : Stochastic Volatility in Trend and Gap Inflation	81
R.59	$\mathcal{M}_2$ : Trend and Gap Inflation	82
R.60	$\mathcal{M}_2$ : Static Parameters	83
R.61	$\mathcal{M}_2$ : Stochastic Volatility in Trend and Gap Inflation	84
R.62	$\mathcal{M}_2$ : Time-Varying SI Parameter	85
R.63	$\mathcal{M}_2$ : Time-Varying AR Coefficient in the Inflation Gap Process	86
R.64	$\mathcal{M}_3$ : Trend and Gap Inflation	87
R.65	$\mathcal{M}_3$ : Static Parameters	88
R.66	$\mathcal{M}_3$ : Stochastic Volatility in Trend and Gap Inflation	89
R.67	$\mathcal{M}_3$ : Time-Varying AR Coefficient in the Inflation Gap Process	90
R.68	$\mathcal{M}_0$ (w/o noise): Trend and Gap Inflation	93
R.69	$\mathcal{M}_0$ (w/o noise): Static Parameters	94
R.70	$\mathcal{M}_0$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation	95
R.71	$\mathcal{M}_0$ (w/o noise): Time-Varying SI Parameter	96
R.72	$\mathcal{M}_1$ (w/o noise): Trend and Gap Inflation	97
R.73	$\mathcal{M}_1$ (w/o noise): Static Parameters	98
R.74	$\mathcal{M}_1$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation	99
R.75	$\mathcal{M}_2$ (w/o noise): Trend and Gap Inflation	100
R.76	$\mathcal{M}_2$ (w/o noise): Static Parameters	101
R.77	$\mathcal{M}_2$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation	102
R.78	$\mathcal{M}_2$ (w/o noise): Time-Varying SI Parameter	103



R.79	$\mathcal{M}_2$ (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process	104
R.80	$\mathcal{M}_3$ (w/o noise): Trend and Gap Inflation	105
R.81	$\mathcal{M}_3$ (w/o noise): Static Parameters	106
R.82	$\mathcal{M}_3$ (w/o noise): Stochastic Volatility in Trend and Gap Inflation	107
R.83	$\mathcal{M}_3$ (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process	108
R.84	log MDD for different numbers of particles (CPI)	110
R.85	log MDD for different numbers of particles (w/o noise)	111
R.86	Relative ESS	113
R.87	Relative ESS (w/o noise)	114
R.88	Simulated Bias in Particle Learning Estimates of Trend and Gap Inflation	118
R.89	Simulated Bias in Particle Learning Estimates of Stochastic Volatilities	119
R.90	Simulated Bias in Particle Learning Estimates $\theta_t$ and $\lambda_t$	120
R.91	Simulated Bias in Particle Learning Estimates of Static Parameters	121
R.92	Detection Rates of True Model	123
R.93	Detection Rates against $\mathcal{M}_1$ Model	123
R.94	Detection Rates of True Model when Measurement Error is Small	124
R.95	Detection Rates against $\mathcal{M}_1$ Model when Measurement Error is Small	124

## R.1 Model-based inflation forecasts and the SPF

Our state space models extract estimate forecasts that reflect rational expectations (RE) and sticky information (SI) forecast updating. The model estimates are conditioned on data of observed SPF forecasts and realized inflation. This section of the results appendix compares forecasts of the average SPF respondent against model-implied RE and SI forecasts.

Section R.1.1 provides a visual comparison of the term structures of SPF forecasts, and model-based RE and SI forecasts at different points in time. Section R.1.2 reports results from forecast comparison between SPF and model-based RE and SI forecasts.

### R.1.1 SPF term structure of inflation forecasts

Here we report briefly on variations in the term structure of inflation forecasts over our data sample ranging from 1968:Q2 to 2018:Q3. First, the raw data is displayed in figure R.1, which shows SPF nowcast and forecasts (omitting, for brevity, data for the two-quarter ahead forecast), and realized inflation. (The figure complements figure 1 in the paper, where panels (a) and (d) are shown.)

Second, we compare observed SPF forecasts with the term structures of RE and SI forecasts in figure R.2. For three selected points in time, the figure shows model based RE and SI forecasts of inflation  $h$  steps ahead, that is  $\mathbf{E}_t \pi_{t+h}$  and  $\mathbf{F}_t \pi_{t+h}$  for  $h = 1, 2, \dots, 5$  as well as the corresponding data points of SPF forecasts. By construction, the SPF  $h$ -quarter ahead forecast is identical to the SI forecast,  $\mathbf{F}_t \pi_{t+h}$ , plus measurement error.<sup>1</sup> Model-based estimates are generated by the  $\mathcal{M}_2$  model, that is most preferred by the data.<sup>2</sup>

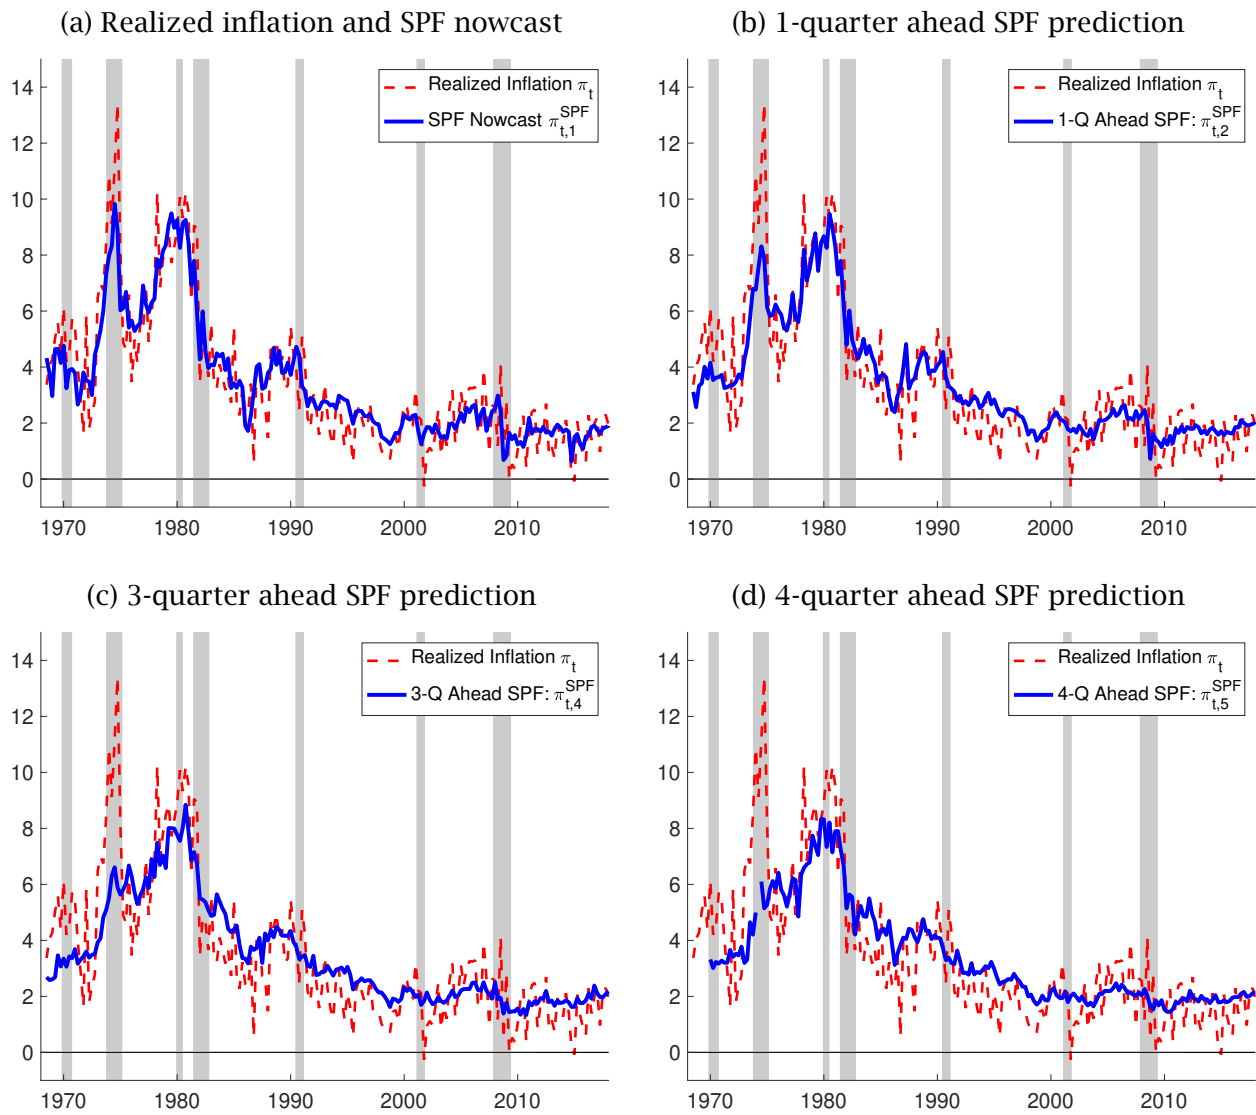
The three snapshots shown in figure R.2 correspond to points in time where stickiness was particularly high or low: 1974:Q4, when stickiness was as low as  $\lambda_{t|T} = 0.19$ , 1996:Q4, by when stickiness had drifted up to  $\lambda_{t|T} = 0.75$ , and 2017:Q4, when stickiness hovered at  $\lambda_{t|T} = 0.62$ .<sup>3</sup> The first snapshot falls into the stagflation recession of 1974, the second into the opportunistic disinflation of the 1990s, and the third into the normalization period after the global financial crisis of 2008/09. As a result, the term structures of inflation forecasts differ quite a bit across the three panels of figure R.2. In 1974:Q4, observed SPF forecasts, as well as model-implied RE and SI predictions, display a downward sloping term structure, as much of the preceding rise in inflation is attributed to a persistent, but not permanent, rise in gap inflation. In contrast, panels (b) and (c) of the figure correspond to relatively stable periods with relatively flat term structures of inflation forecasts. However, as seen in panel (b), compared to the RE case, the level of the SI term structure runs somewhat higher during the period of opportunistic disinflation. As discussed in the paper, the SI trend has been slower to

<sup>1</sup>Recall our timing protocol, whereby SPF forecasts collected in the middle of quarter  $t + 1$  are formed conditional on information available at the end of quarter  $t$ . The SPF nowcast collected in the middle of quarter of  $t + 1$  thus maps into  $\mathbf{F}_t \pi_{t+1}$  plus noise, and so on for other values of  $h$ .

<sup>2</sup>In model  $\mathcal{M}_2$ , the sticky-information weight,  $\lambda_t$ , is drifting, and so is the inflation gap persistence coefficient,  $\theta_t$ .

<sup>3</sup>The estimated values of  $\lambda_t$  quoted, refer to smoothed estimates obtained from the  $\mathcal{M}_2$  model as well.

Figure R.1: Realized Inflation and SPF Predictions



Note: Data from 1968Q4 to 2018Q3. Vertical gray bands denote NBER dated recessions.

recognize the ongoing disinflation efforts during that period.<sup>4</sup> Toward the end of our sample, as illustrated in panel (c), differences between SI and RE term structure are negligible.

Strikingly, despite marked difference in underlying forecast stickiness, there seems to be much less difference *in differences* between the term structures of RE and SI forecasts when compared across the panels of figure R.2. Apart from the difference in levels in the 1990s discussed above, SI and RE terms structures are fairly alike, regardless of whether  $\lambda_t$  is low (as in the 1970s) or high (as by 1996 or the end of the sample). This broad similarity in term structures across variations in  $\lambda_t$  confirms our narrative that variations in  $\lambda_t$  may be state-dependent, and that high stickiness might occur only at times, when a slower frequency of updating SI forecasts is not costly. We pick up this theme in the next section, which provides a forecast comparison between SI and RE forecasts.

### R.1.2 Forecasting Inflation with Our State Space Models

This appendix reports on a small forecasting exercise. We compare the  $h$ -quarter ahead inflation forecasts produced by  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$ , both under rational expectations (RE) and sticky information (SI), to the average SPF inflation predictions of the same horizon. The forecast comparison is applied to our full data sample, from 1968Q4 to 2018Q3 sample, as well as for two sub-samples that cover the pre and post Great Moderation periods, respectively. Table R.1 displays displays results for our baseline case with noise in the inflation process. For the case of no noise in the inflation process, similar results are presented in table R.2.

The model-based forecasts used for the forecast comparison reflect *filtered* estimates of trend and persistent gap inflation generated from the particle learning filter. A critical benefit of the particle learning estimates of model-implied inflation forecasts is that the underlying parameters are also filtered, and reflect only sample information up to the time the forecast is made. In short, the setup corresponds to the use of a *growing* estimation window. For all sub-samples considered, forecasts reflect particle learning estimates that condition on information available since the beginning of our full data sample as indicated in each table. The sub-sample comparison thus serves only to break down over which periods the full-sample forecast errors accrued, the sub-sample analysis does not consider variations in forecasts induced by different starting points of the estimation window.


Table R.1 reports the root mean square errors (RMSEs) of the  $h$ -quarter ahead SPF inflation forecasts and relative RMSEs of model-based RE and SI forecasts that use the SPF predictions as benchmark. As discussed in the main paper, SPF responses collected in the middle of quarter  $t$  are treated as forecasts made at the end of quarter  $t - 1$  by our models, so that the forecast horizon of the SPF nowcast is labelled  $h = 1$ , the one-quarter ahead SPF predictions is labelled  $h = 2$  and so on. This timing assumption provides the SPF some informational advantage, in particular for near-term forecasts, as SPF respondents have access to almost half-a-quarter's worth of extra information. Indeed,

---


<sup>4</sup>See also Mertens (2016).

Figure R.2: Term structures of SPF data, SI and RE forecasts


(a) 1974:Q4 ( $\lambda_{t|T} = 0.19$ )



(b) 1996:Q4 ( $\lambda_{t|T} = 0.75$ )



(c) 2017:Q4 ( $\lambda_{t|T} = 0.62$ )



Note: Rational expectations (RE) and sticky information (SI) forecasts as generated from filtered estimates of the  $\mathcal{M}_2$  model that is preferred by the data.

across all comparison tables, SPF forecasts outperform their model-based alternatives by up to 10% in terms of relative RMSE (which correspond to about 10 basis points in terms of RMSE levels).

The timing effect is most visible in the comparison of SPF forecasts against model-based SI forecasts. The latter differ from SPF data only by the measurement errors imputed by each model, which by themselves should not lead to marked differences in forecasting. Indeed, the RMSE of SPF and model-based SI forecasts are nearly indistinguishable for horizons greater than one, when the informational timing advantage begins to dissipate.

A more interesting conclusion emerges from comparing model-based RE forecasts against the SPF at horizons  $h > 1$ . As shown in table R.1, our models generally perform at par or slightly better than the SPF; however, the relative RMSE are typically just around 0.98 and drop to 0.95 at best. If anything, models  $\mathcal{M}_0$  and  $\mathcal{M}_2$  where the sticky-information weight  $\lambda_t$  is time-varying, record slightly larger improvements in relative RMSE than the other two. Interestingly, any outperformance of the model-based RE forecasts accrues mostly prior to the year 2000, when inflation had stabilized.<sup>5</sup> As discussed in the main paper, our estimates also suggest that the sticky-information weight  $\lambda_t$  has risen during the post-Volcker period to levels indicating a relatively low frequency of forecast updating by SPF respondents (or, in other words, a high level of stickiness). Taken together, the relatively slim forecasting margin of RE forecasts, in particular over the period when  $\lambda_t$  has been high,<sup>6</sup> suggest that variations in forecast stickiness occurred not independently of their likely costs in terms of forecasters RMSEs, since even periods of high stickiness seem to have coincided with only small losses (if any) in the performance of SPF and SI forecasts.

---

<sup>5</sup>By the year 2000, not only had the Volcker disinflation been over, but the efforts of opportunistic disinflation of the 1990s had also run their course and trend inflation had settled around 2%.

<sup>6</sup>By construction, when stickiness is estimated to be low, SI forecasts (and thus also the SPF) are close to RE.

Table R.1: Inflation Forecast Comparison

horizon	SPF (RMSE)	RE model forecasts (rel. RMSE)				SI model forecasts (rel. RMSE)			
		$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
PANEL (A): 1968:Q3 to 2018:Q3									
1	1.09	1.13	1.03	1.09	1.10	1.11	1.05	1.09	1.11
2	1.40	1.01	1.00	1.00	1.02	1.02	1.00	1.01	1.02
3	1.58	0.98	0.99	0.98	1.00	1.00	1.00	1.00	1.01
4	1.71	0.96	0.99	0.98	0.99	1.00	1.00	1.00	1.00
5	1.84	0.98	0.99	0.98	0.99	0.99	1.00	1.00	0.99
PANEL (B): 1968:Q3 to 1984:Q4									
1	1.48	1.15	1.03	1.09	1.10	1.11	1.06	1.09	1.12
2	2.04	1.00	1.00	1.00	1.01	1.01	1.00	1.01	1.01
3	2.37	0.98	0.98	0.98	0.99	1.00	0.99	1.00	1.00
4	2.58	0.95	0.98	0.97	0.98	0.99	0.99	0.99	0.99
5	2.85	0.95	0.96	0.96	0.96	0.97	0.97	0.97	0.96
PANEL (C): 1985:Q1 to 2000:Q1									
1	0.88	1.07	1.01	1.05	1.10	1.10	1.03	1.09	1.10
2	1.01	1.02	0.99	0.99	1.03	1.03	1.00	1.02	1.03
3	1.09	0.96	0.99	0.96	1.00	1.01	1.00	1.01	1.01
4	1.13	0.96	1.01	0.98	1.01	1.03	1.02	1.03	1.02
5	1.17	0.97	1.00	0.98	1.00	1.02	1.01	1.02	1.01
PANEL (D): 2000:Q1 to 2018:Q3									
1	0.79	1.12	1.05	1.11	1.11	1.10	1.05	1.10	1.10
2	0.85	1.03	1.03	1.02	1.02	1.02	1.03	1.02	1.02
3	0.87	1.02	1.02	1.01	1.03	1.03	1.02	1.03	1.03
4	0.90	1.03	1.03	1.03	1.04	1.05	1.04	1.05	1.05
5	0.96	1.01	1.00	1.01	1.01	1.01	1.00	1.01	1.01

Note: RMSE of SPF forecasts and relative RMSE of RE and SI predictions generated by a given model compared to the SPF forecast. (Numbers below one indicate a lower RMSE of the model forecasts). In each panel, model forecasts used are based on filtered estimates using data since 1968:Q3. Forecast errors are then collected over the (sub)periods indicated in each panel.

Table R.2: Inflation Forecast Comparison (models with noise in inflation)

horizon	SPF (RMSE)	RE model forecasts (rel. RMSE)				SI model forecasts (rel. RMSE)			
		$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
PANEL (A): 1968:Q3 to 2018:Q3									
1	1.09	1.19	1.11	1.13	1.12	1.13	1.09	1.12	1.12
2	1.40	1.03	1.02	1.03	1.02	1.02	1.01	1.02	1.02
3	1.58	0.99	0.99	0.99	0.99	1.01	1.00	1.00	1.00
4	1.71	0.96	0.99	0.98	0.99	1.00	1.00	1.00	1.00
5	1.84	0.97	0.98	0.98	0.98	0.99	0.99	0.99	0.99
PANEL (B): 1968:Q3 to 1984:Q4									
1	1.48	1.18	1.11	1.12	1.12	1.13	1.09	1.11	1.13
2	2.04	1.02	1.01	1.02	1.01	1.02	1.01	1.01	1.02
3	2.37	0.99	0.99	1.00	0.99	1.00	1.00	1.00	1.00
4	2.58	0.96	0.98	0.99	0.98	0.99	0.99	0.99	0.99
5	2.85	0.95	0.96	0.96	0.96	0.97	0.96	0.97	0.96
PANEL (C): 1985:Q1 to 2000:Q1									
1	0.88	1.15	1.08	1.16	1.14	1.15	1.08	1.14	1.13
2	1.01	1.05	1.02	1.06	1.04	1.05	1.03	1.04	1.04
3	1.09	0.97	1.00	0.95	1.00	1.01	1.01	1.00	1.00
4	1.13	0.95	1.01	0.94	1.00	1.01	1.02	1.01	1.01
5	1.17	0.94	0.98	0.92	0.98	0.99	0.99	0.98	0.99
PANEL (D): 2000:Q1 to 2018:Q3									
1	0.79	1.25	1.12	1.15	1.11	1.12	1.10	1.11	1.10
2	0.85	1.06	1.05	1.06	1.04	1.04	1.04	1.03	1.04
3	0.87	1.03	1.02	1.01	1.03	1.04	1.02	1.03	1.03
4	0.90	1.03	1.03	1.02	1.04	1.05	1.04	1.05	1.05
5	0.96	1.01	1.02	1.01	1.01	1.01	1.02	1.01	1.01

Note: RMSE of SPF forecasts and relative RMSE of RE and SI predictions generated by a given model compared to the SPF forecast. (Numbers below one indicate a lower RMSE of the model forecasts). In each panel, model forecasts used are based on filtered estimates using data since 1968:Q3. Forecast errors are then collected over the (sub)periods indicated in each panel.



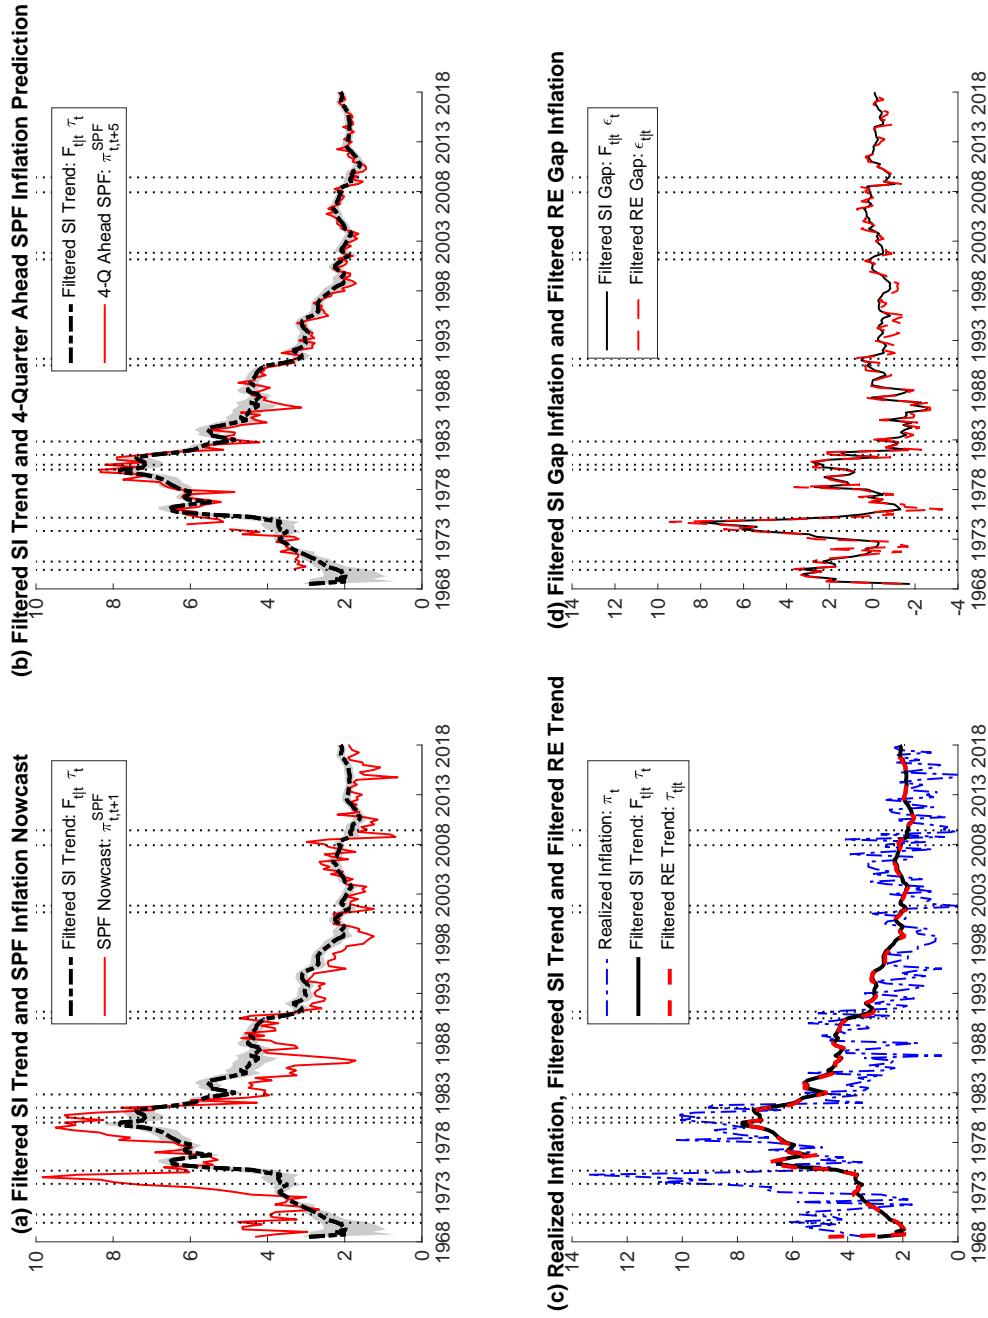
## R.2 Detailed state-space model results

This section reports detailed estimates for the various model parameters and latent states of our four model variants  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$ . First, section R.2.1 shows results where, in each model, the inflation process contains a serially-uncorrelated irregular component, as in equations (1) and (2) of the paper. (We refer to this irregular component also as noise in the inflation process.) Section R.2.2 then turns to results based on an alternative specification for inflation that omits the irregular component. Sections R.2.3, and R.2.4 report log MDDs over time and for different numbers of particles and effective sample sizes.

### R.2.1 Results for the four model variants

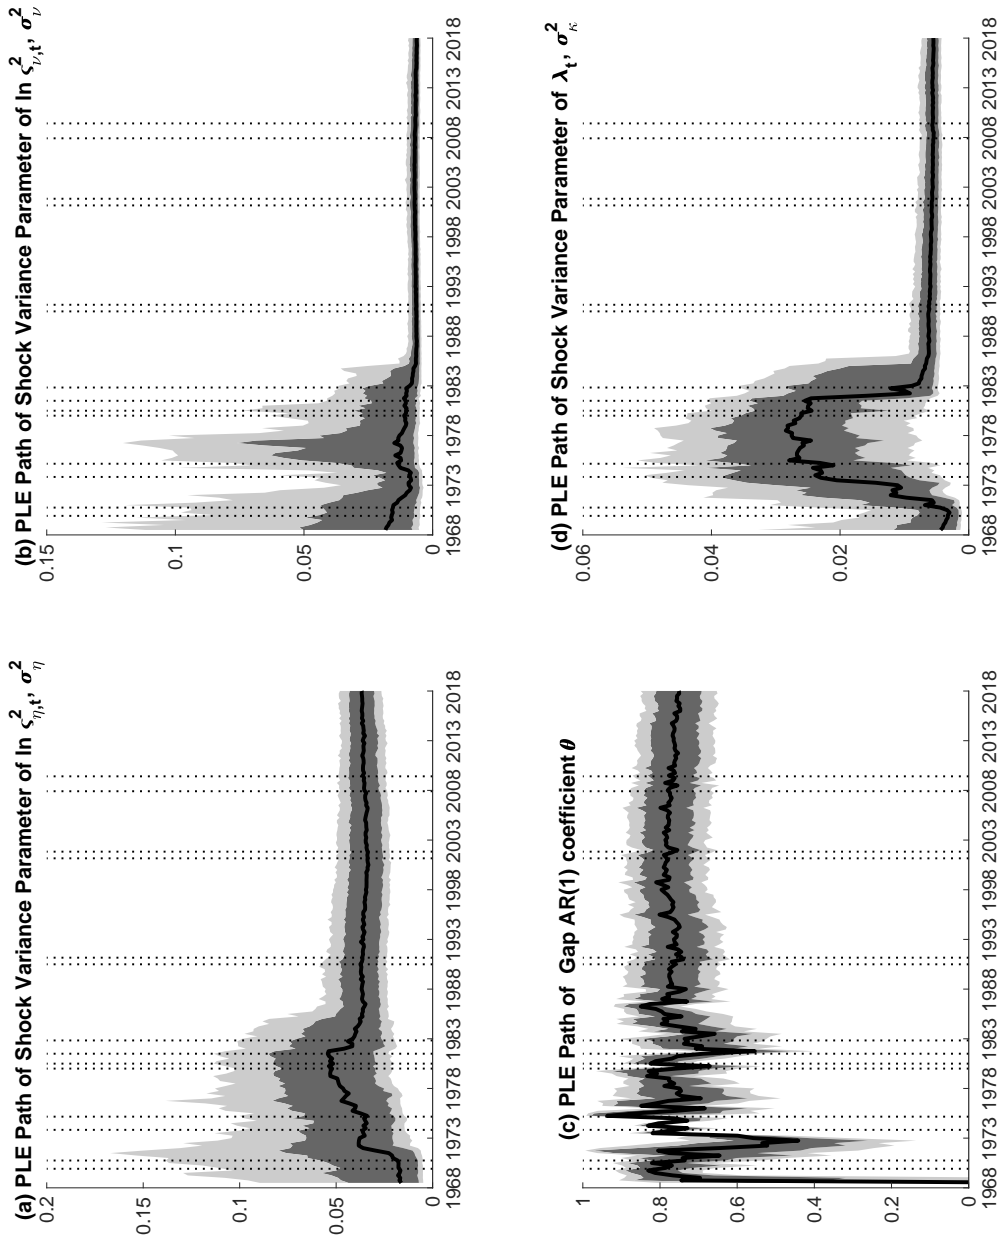
Figures R.3-R.6 report results for  $\mathcal{M}_0$ , with time-varying  $\lambda_t$  but constant  $\theta$ . Figures R.7-R.9 show results for  $\mathcal{M}_1$ , where both  $\lambda$  and  $\theta$  are assumed to be constant. Figures R.10-R.14 provide results for  $\mathcal{M}_2$ , where both  $\lambda_t$  and  $\theta_t$  are time-varying. Finally, Figures R.15-R.18 show results for  $\mathcal{M}_3$ , where  $\theta_t$  is time-varying but  $\lambda$  is assumed constant.

Figure R.3:  $\mathcal{M}_0$ : Trend and Gap Inflation



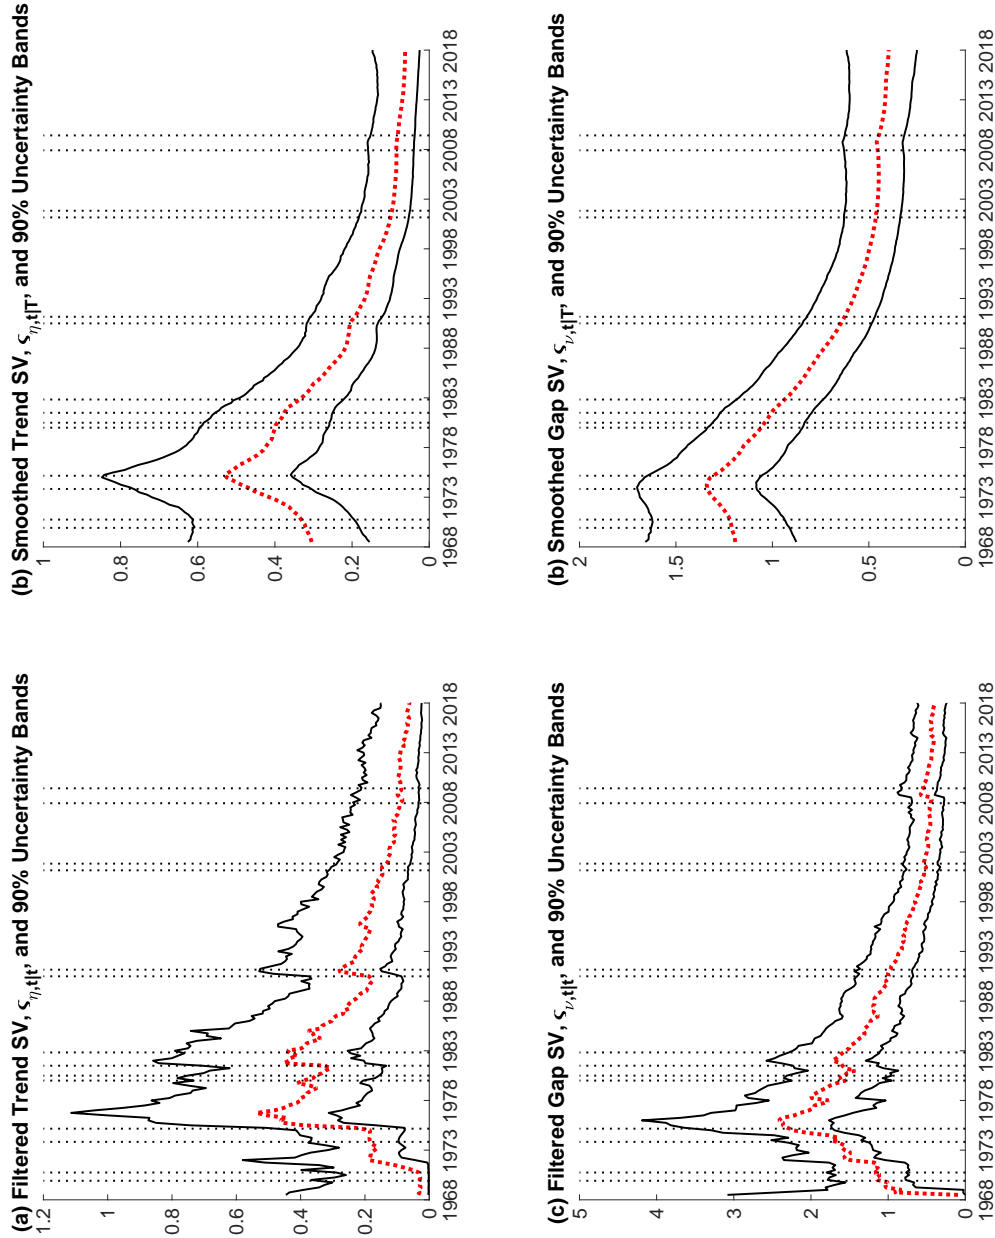
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_0$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.4:  $\mathcal{M}_0$ : Static Parameters




Note: Posterior quantiles of particle-learning estimates. Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_0$ . Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.5:  $\mathcal{M}_0$ : Stochastic Volatility in Trend and Gap Inflation



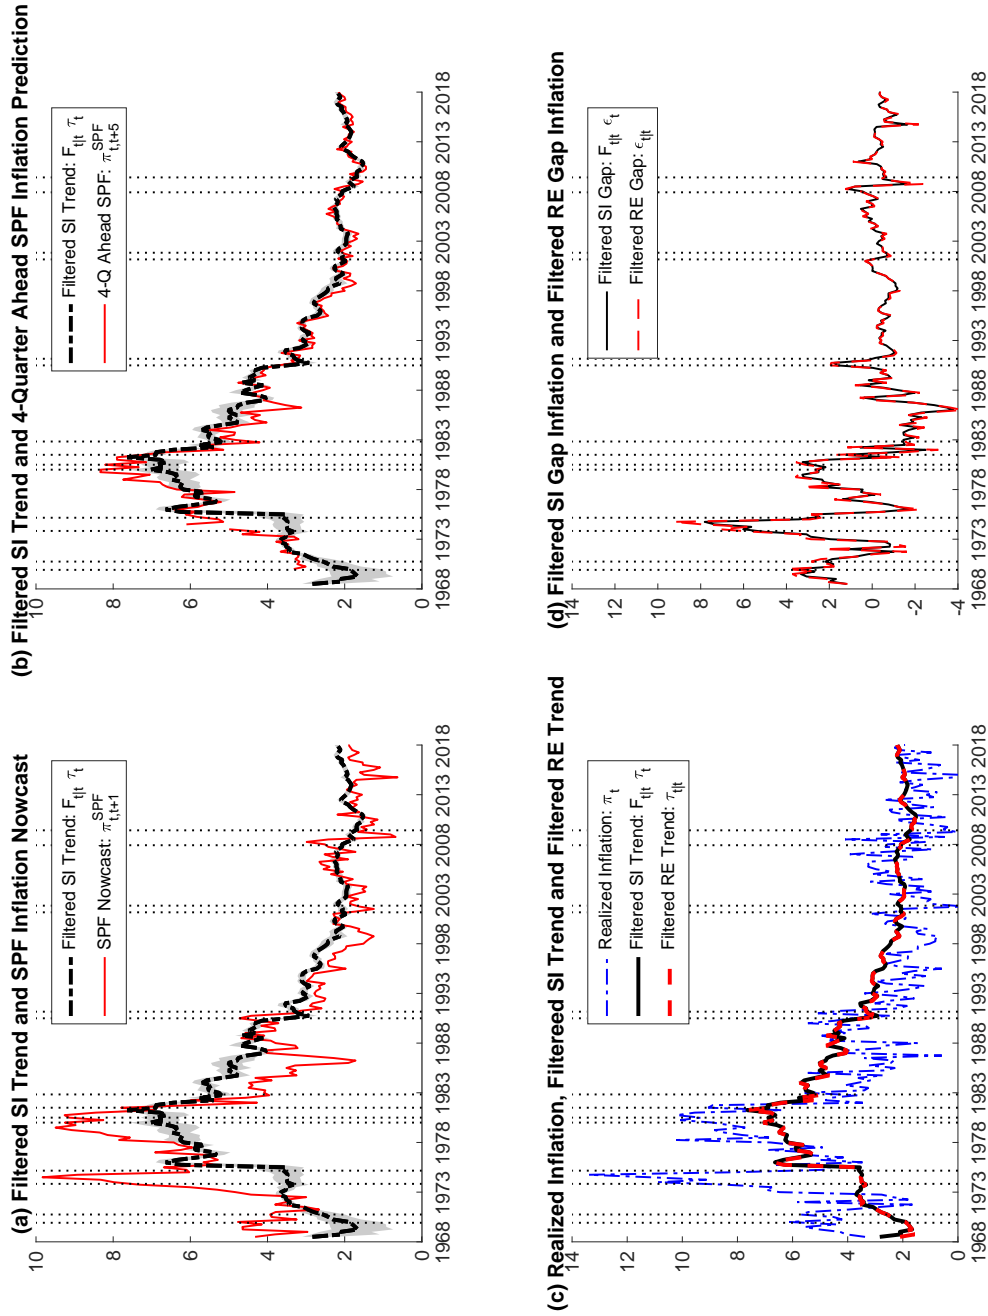
Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_0$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.6:  $\mathcal{M}_0$ : Time-Varying SI Parameter



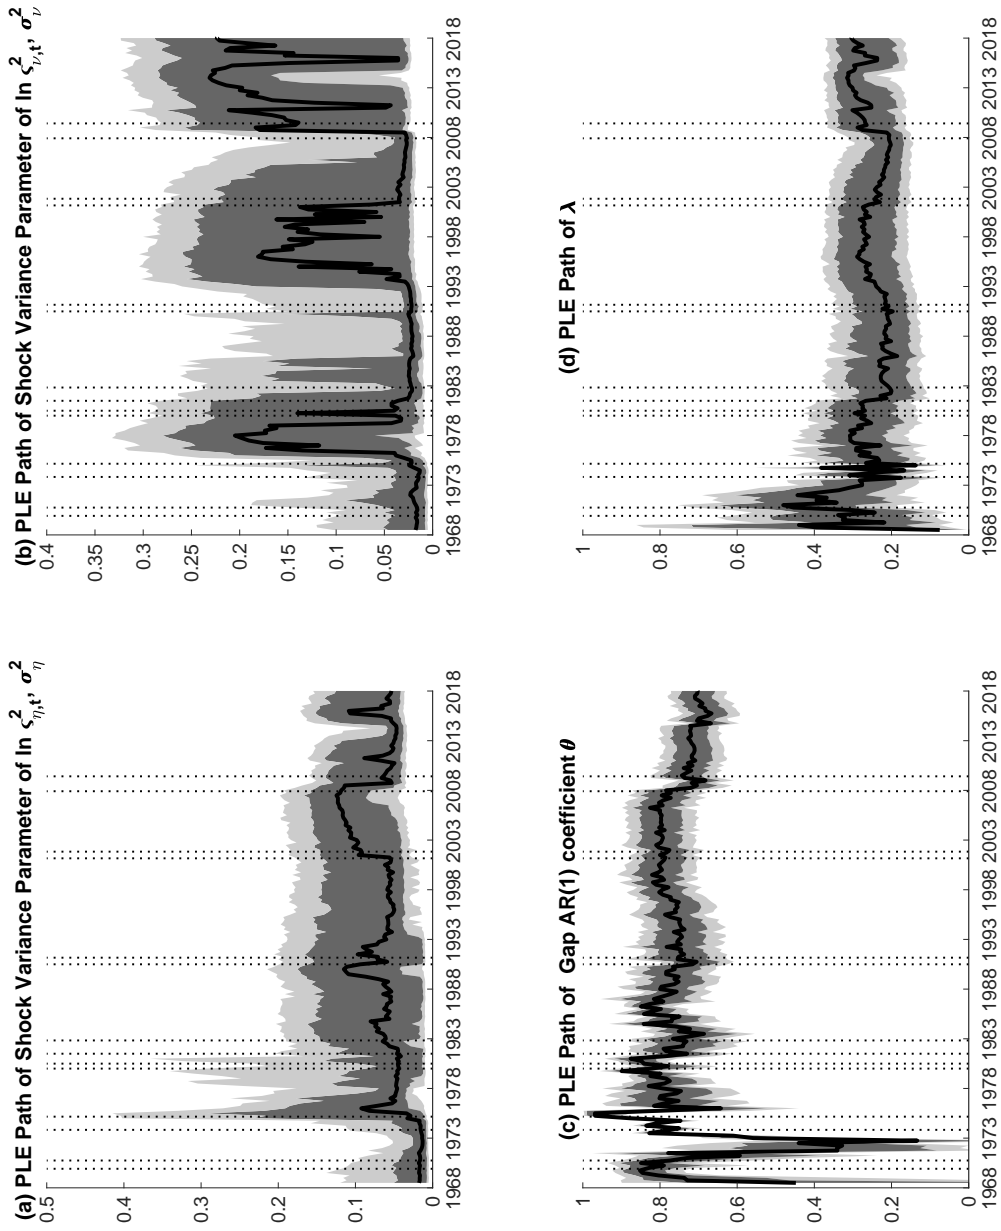
Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{\text{fit}}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{\text{f|T}}$  surrounded by 90% uncertainty bands delineated by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{\text{f|T}} - \lambda_{\text{0|T}}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_0$ .

Figure R.7:  $\mathcal{M}_1$ : Trend and Gap Inflation



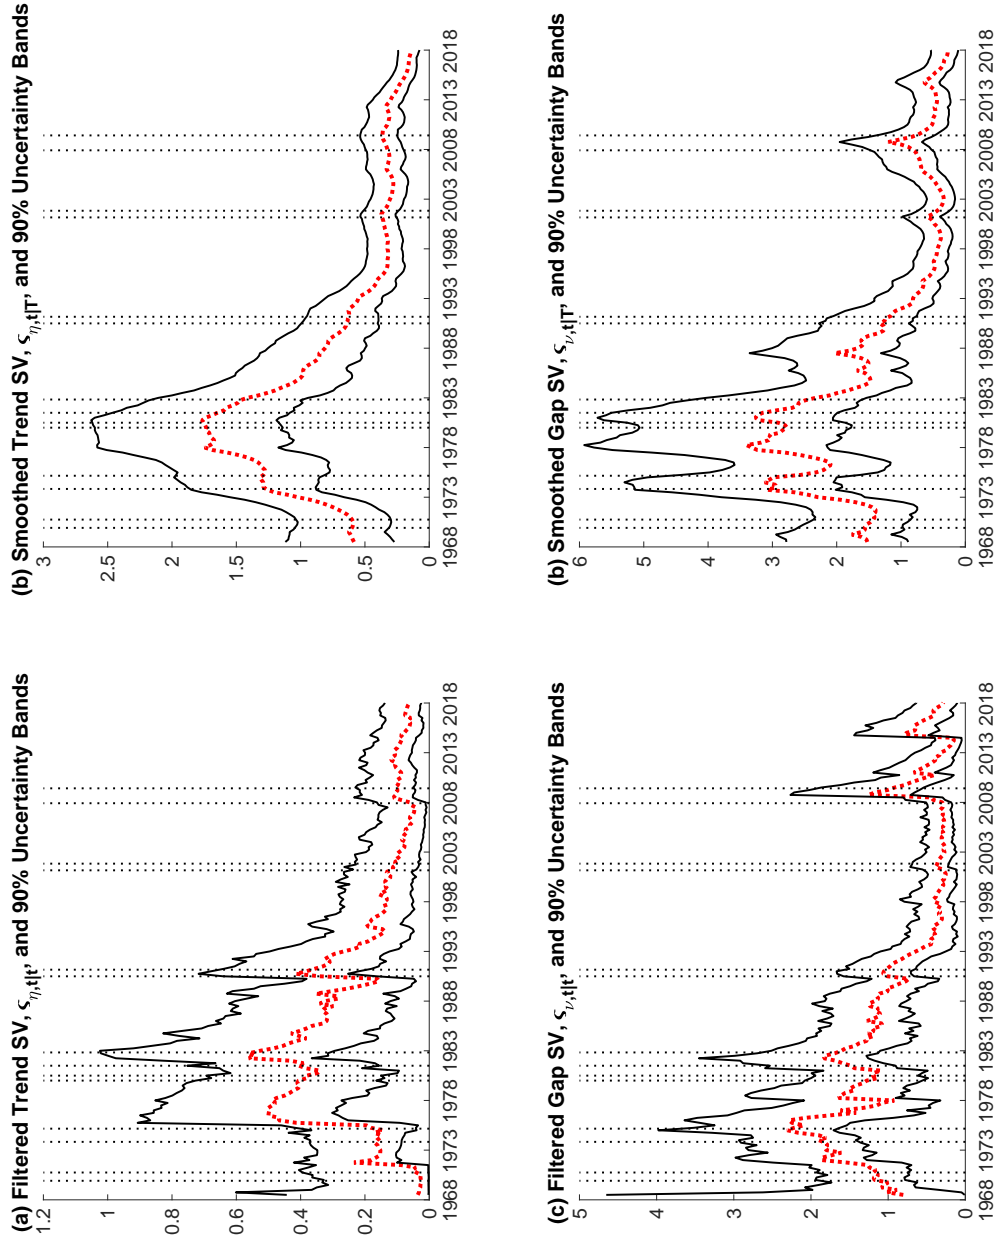
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.8:  $\mathcal{M}_1$ : Static Parameters



Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_1$ . Dotted vertical lines denote NBER recession peaks and troughs.

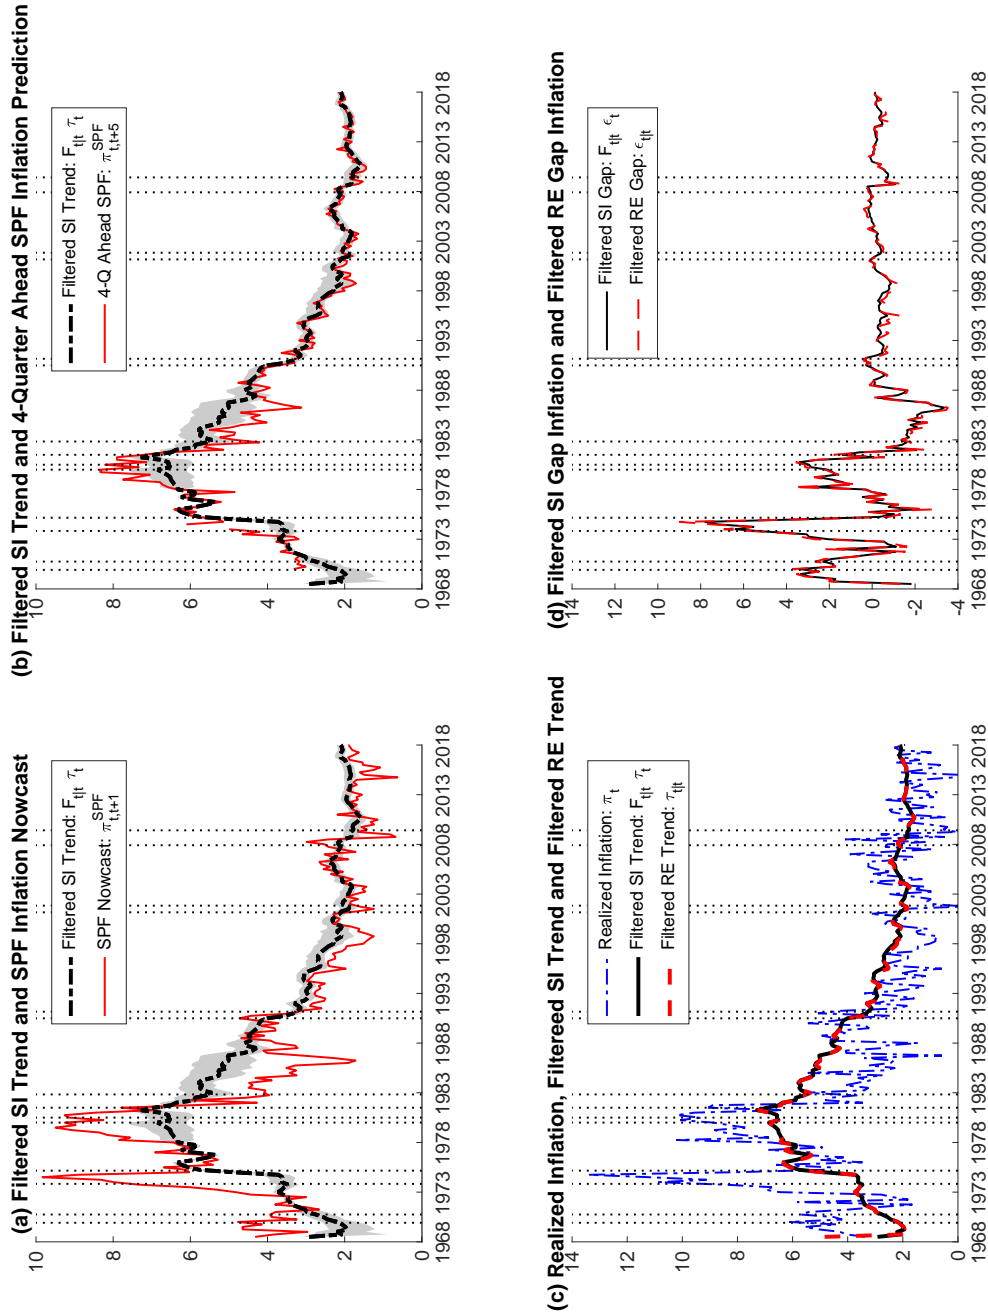
Figure R.9:  $\mathcal{M}_1$ : Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_1$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

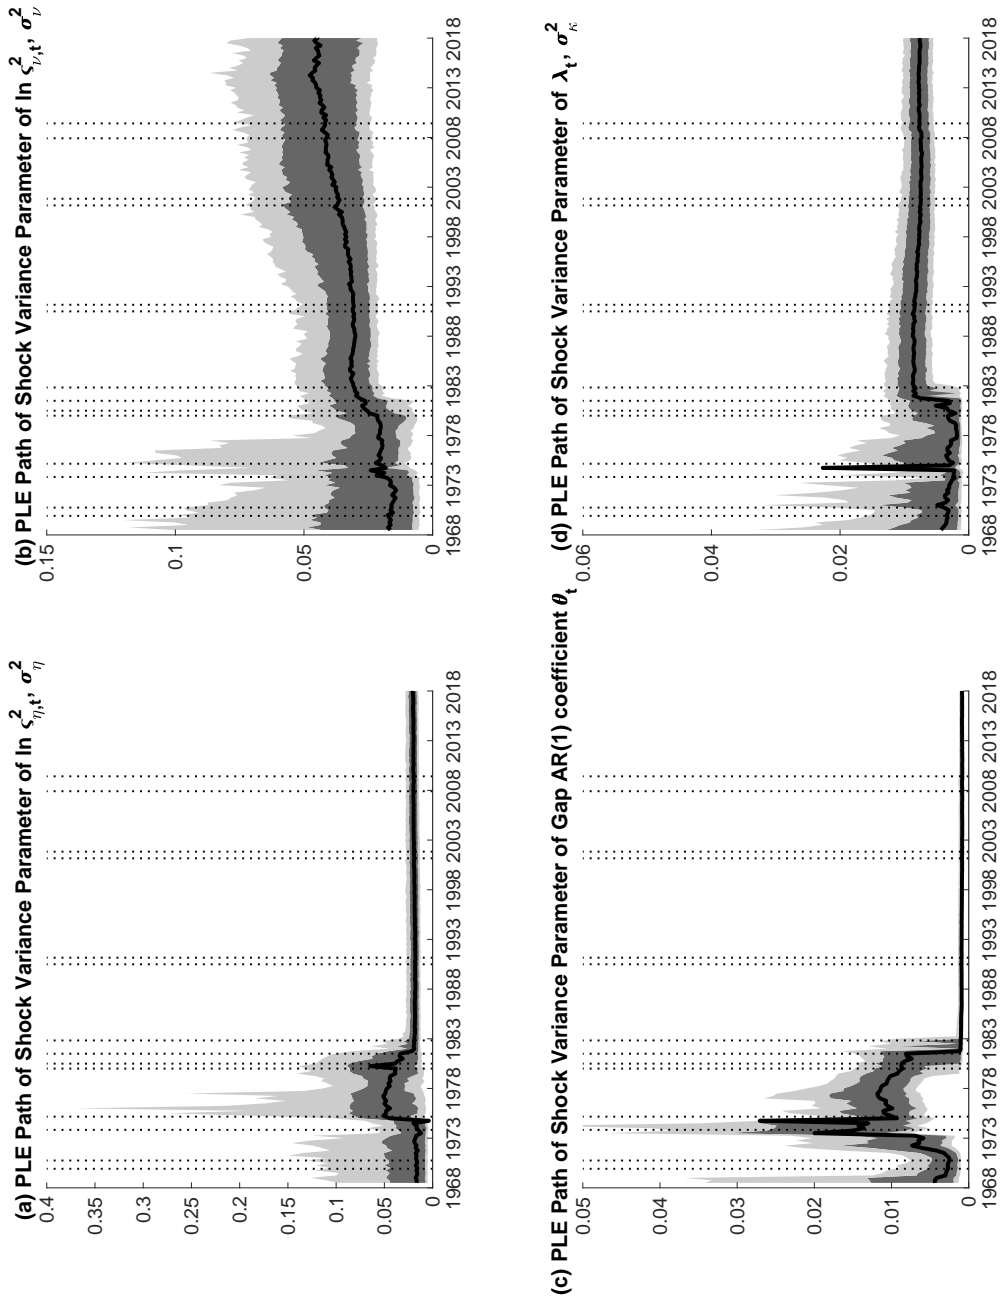


Figure R.10:  $\mathcal{M}_2$ : Trend and Gap Inflation



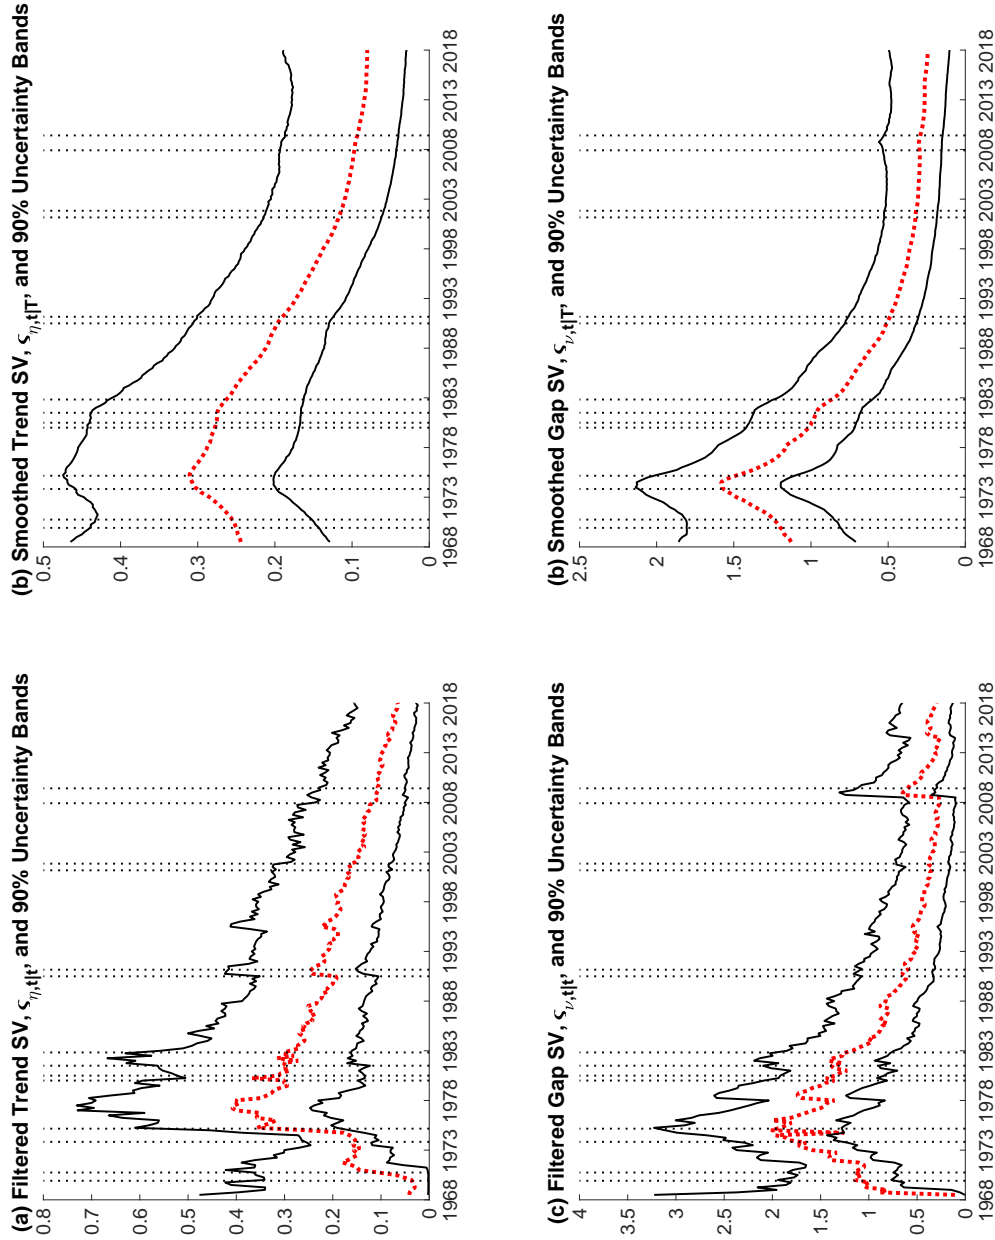
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.11:  $\mathcal{M}_2$ : Static Parameters



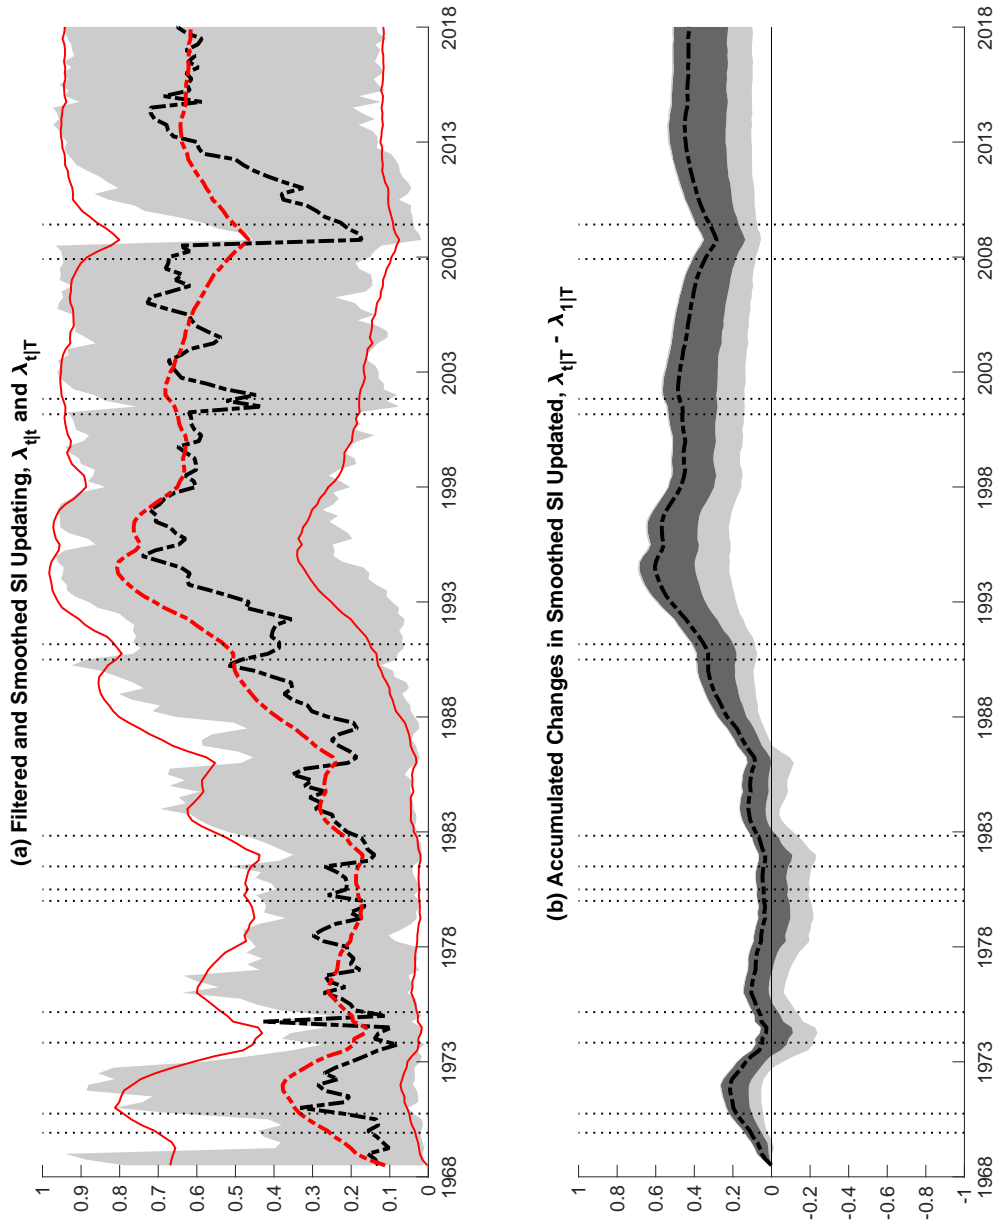
Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_2$ . Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.12:  $\mathcal{M}_2$ : Stochastic Volatility in Trend and Gap Inflation



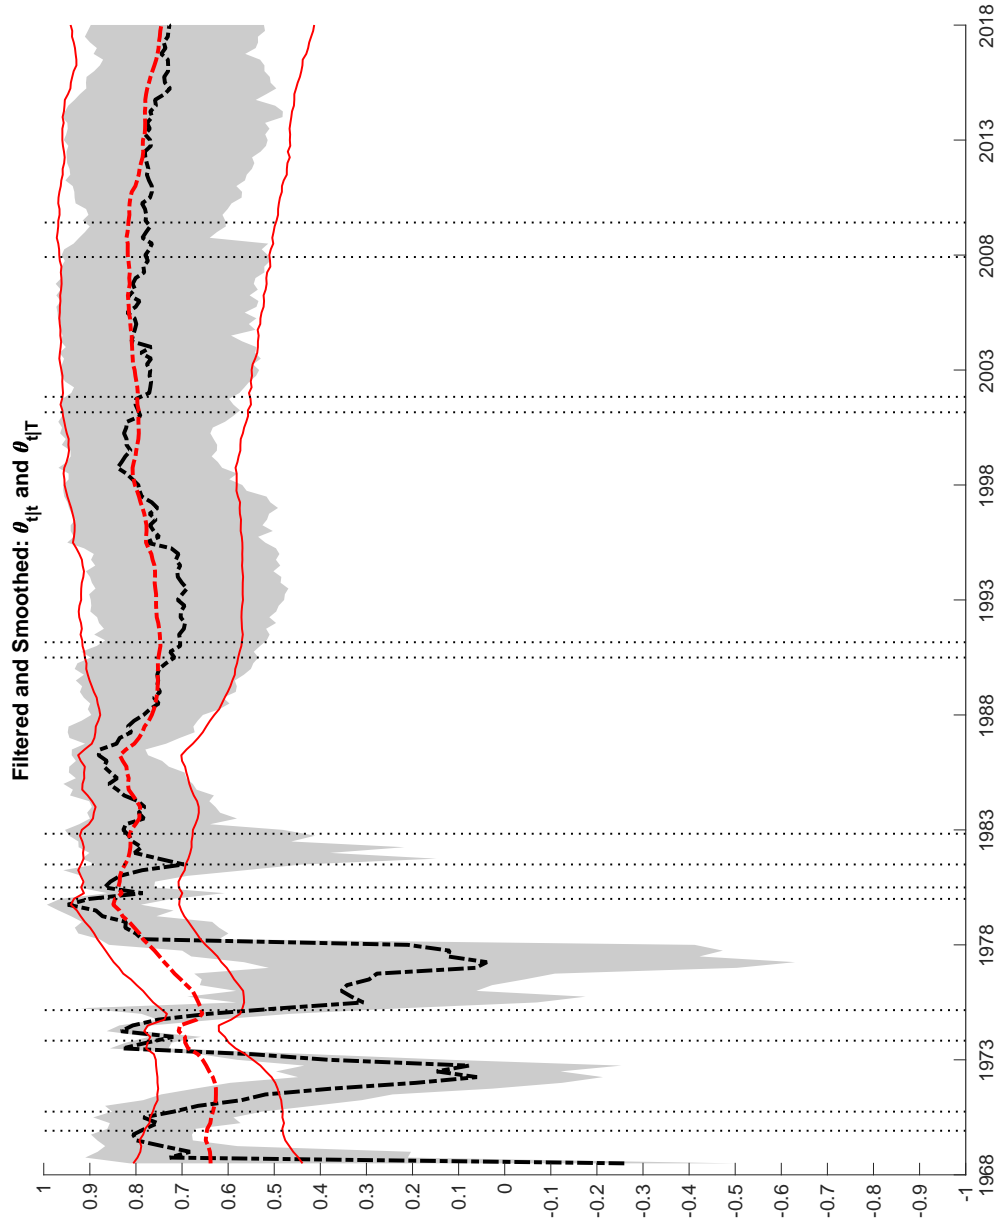
Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_2$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.13:  $\mathcal{M}_2$ : Time-Varying SI Parameter



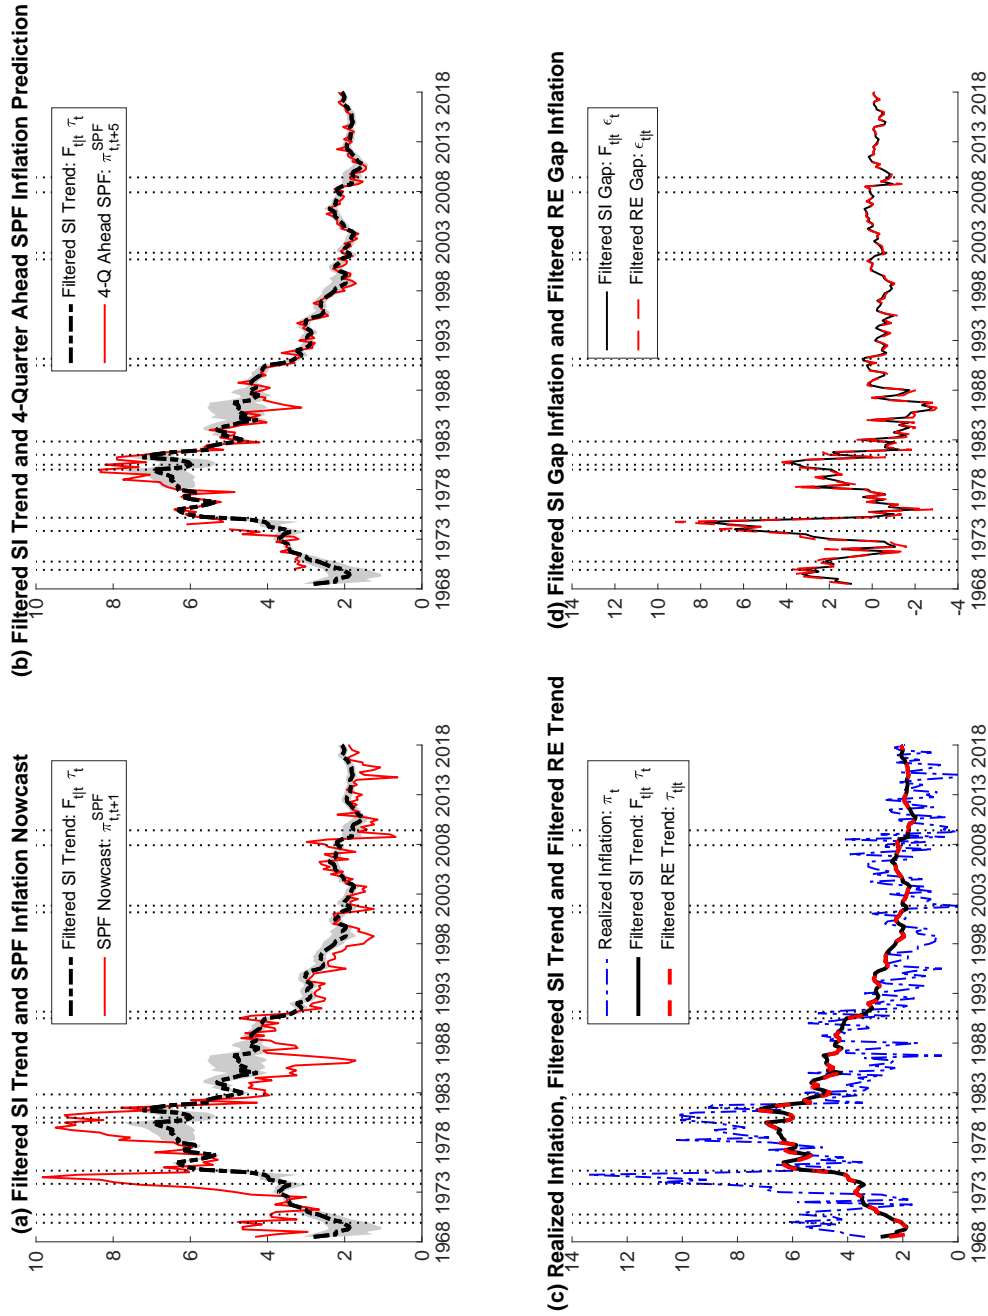
Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{\text{fit}}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{\text{fIT}}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{\text{fIT}} - \lambda_{\text{fIT}}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_2$ .

Figure R.14:  $\mathcal{M}_2$ : Time-Varying AR Coefficient in the Inflation Gap Process



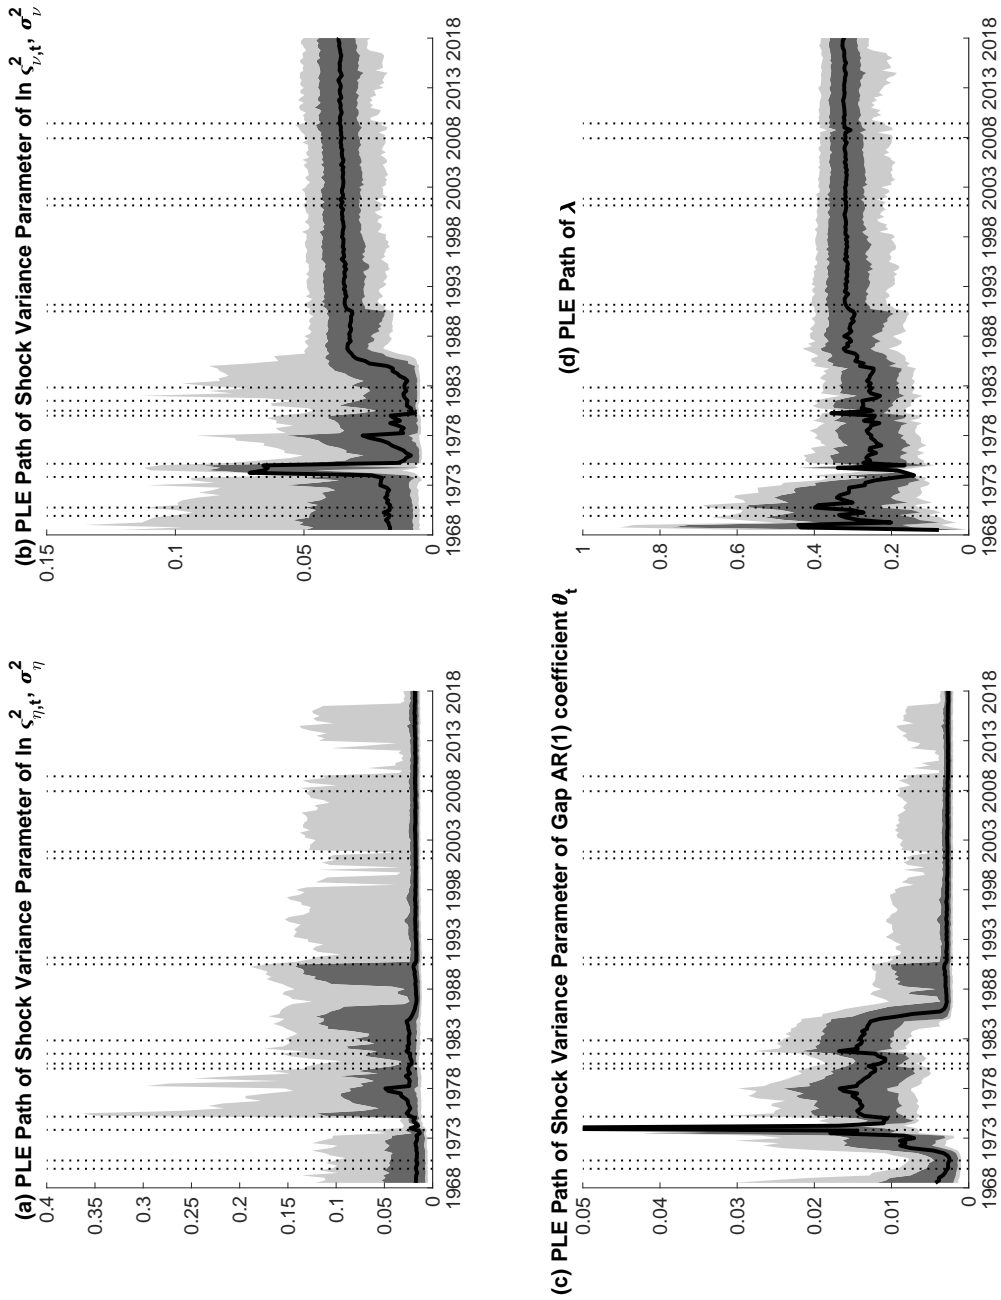
Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{ft}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{ftT}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_2$ .

Figure R.15:  $\mathcal{M}_3$ : Trend and Gap Inflation



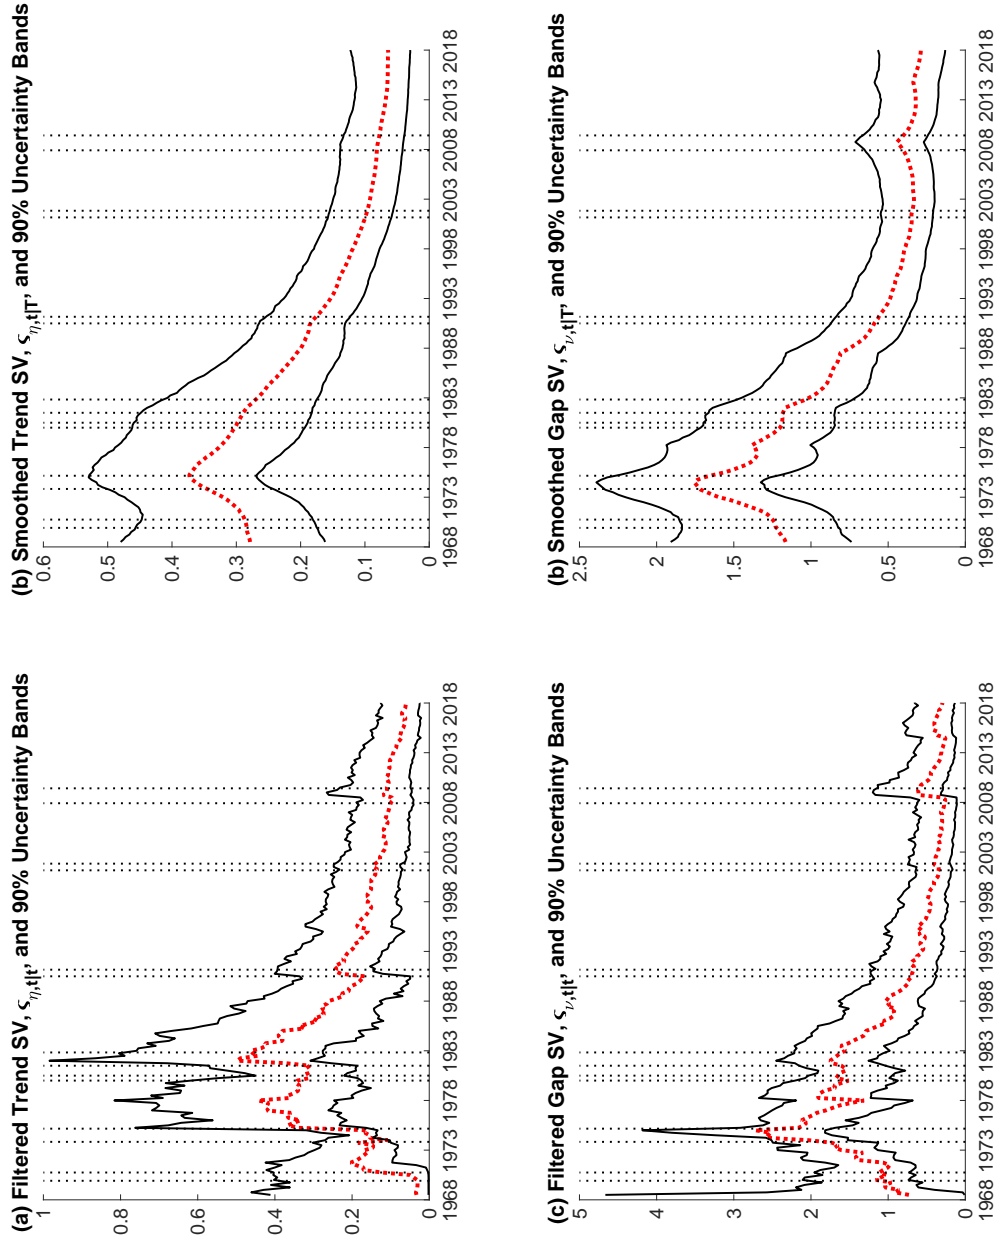
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_3$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.16:  $\mathcal{M}_3$ : Static Parameters



Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_3$ . Dotted vertical lines denote NBER recession peaks and troughs.


Figure R.17:  $\mathcal{M}_3$ : Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_3$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.



Figure R.18:  $\mathcal{M}_3$ : Time-Varying AR Coefficient in the Inflation Gap Process



Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{fit}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{fitT}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_3$ .

## R.2.2 Results for case without noise component in inflation

This section reports results for alternative versions of each of our four model variants. These alternative state space models shut off the irregular component (or noise) in the inflation equation. In this case, the inflation process reduces to

$$\pi_t = \tau_t + \varepsilon_t, \tag{R.1}$$

with the remaining elements of each model, including the specification of priors, unchanged. Please note that, as in the baseline specification, measurement error is retained in the equations mapping SI forecasts into SPF forecast data.

Table R.3 reports estimates of static parameters as well as log MDDs for each of the four model variants when noise in the inflation equation has been removed. In comparison to results from the baseline specification (with noise in inflation), as reported in Table 5 of the paper, the logMDDs are generally lower (and significantly so), providing strong evidence against removing the noise component from the inflation process. However, when comparing these alternative model variants against each other, the  $\mathcal{M}_2$  model with time-varying  $\theta_t$  and  $\lambda_t$  continues to attract the highest log MDD (as reported in the paper for the case with noise in inflation).

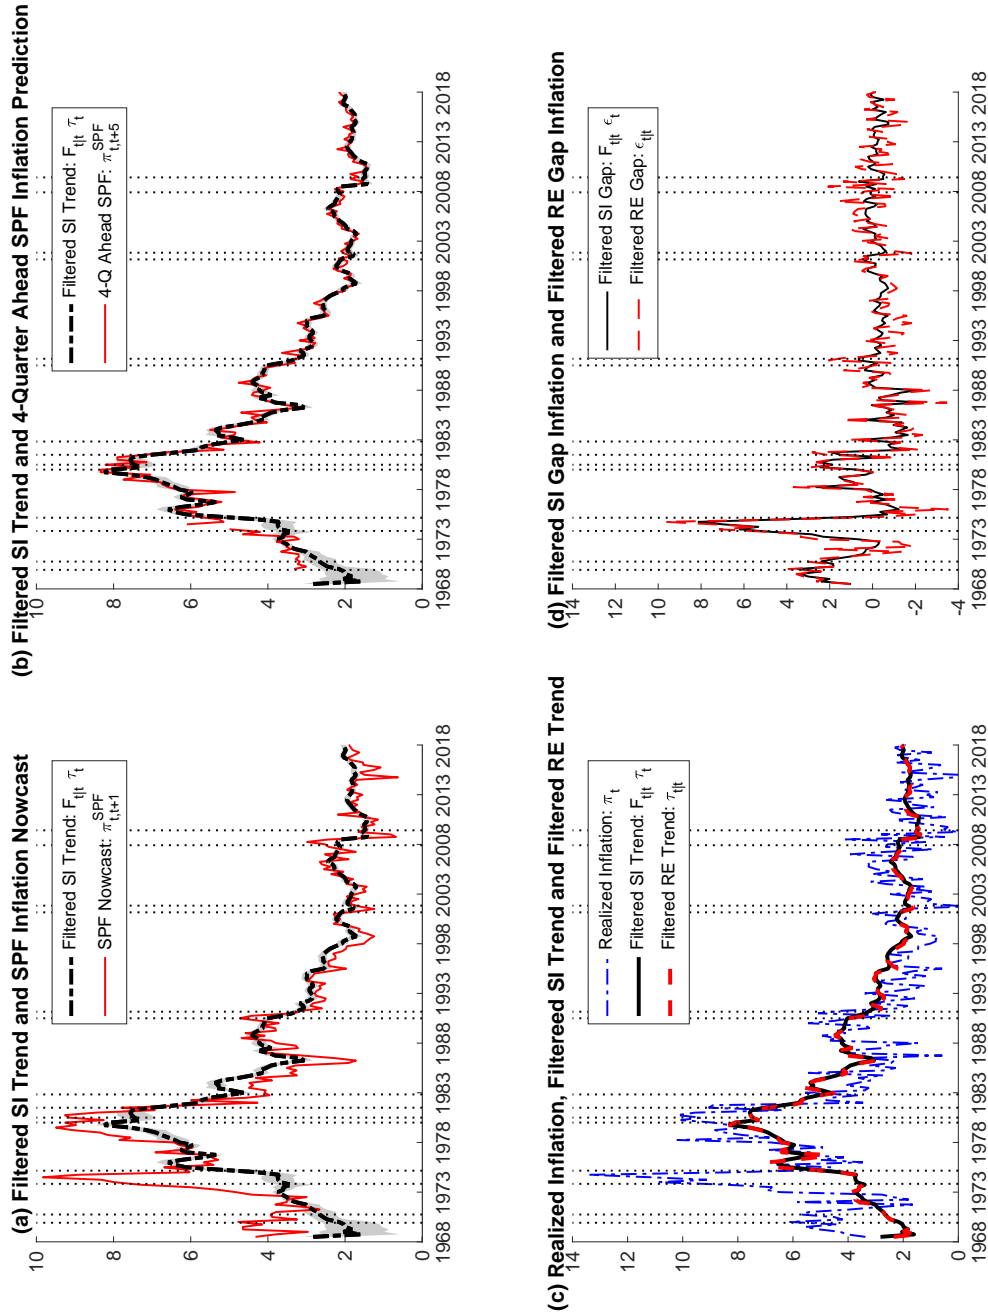
Detailed results are shown separately for each model variant in the following figures: Figures R.19-R.22 report results for  $\mathcal{M}_0$ , with time-varying  $\lambda_t$  but constant  $\theta$ . Figures R.23-R.25 show results for  $\mathcal{M}_1$ , where both  $\lambda$  and  $\theta$  are assumed to be constant. Figures R.26-R.30 provide results for  $\mathcal{M}_2$ , where both  $\lambda_t$  and  $\theta_t$  are time-varying. Finally, Figures R.31-R.34 show results for  $\mathcal{M}_3$ , where  $\theta_t$  is time-varying but  $\lambda$  is assumed constant.

Table R.3: Parameter Estimates and MDDs w/o noise in inflation

Parameter	Models			
	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
Variances of shocks to SV processes				
$\sigma_\eta^2$ (Trend SV)	0.023 [0.014, 0.135]	0.014 [0.012, 0.016]	0.015 [0.011, 0.021]	0.033 [0.021, 0.110]
$\sigma_v^2$ (Gap SV)	0.044 [0.016, 0.065]	0.037 [0.031, 0.045]	0.016 [0.006, 0.025]	0.015 [0.005, 0.025]
Persistence of inflation gap				
$\theta$	0.426 [0.307, 0.544]	0.238 [0.126, 0.355]	-	-
$\sigma_\phi^2$	-	-	0.016 [0.010, 0.024]	0.015 [0.010, 0.026]
Forecast stickiness				
$\lambda$	-	0.235 [0.192, 0.283]	-	0.302 [0.218, 0.368]
$\sigma_\kappa^2$	0.014 [0.009, 0.020]	-	0.002 [0.001, 0.010]	-
Measurement error variances				
$\sigma_{\xi, \pi}^2$	-	-	-	-
$\sigma_{\xi, 1}^2$	0.176 [0.152, 0.208]	0.008 [0.007, 0.010]	0.168 [0.141, 0.202]	0.170 [0.141, 0.202]
$\sigma_{\xi, 2}^2$	0.061 [0.051, 0.073]	0.008 [0.007, 0.010]	0.055 [0.045, 0.068]	0.061 [0.052, 0.073]
$\sigma_{\xi, 3}^2$	0.043 [0.035, 0.052]	0.008 [0.007, 0.010]	0.038 [0.031, 0.050]	0.044 [0.037, 0.052]
$\sigma_{\xi, 4}^2$	0.050 [0.040, 0.061]	0.008 [0.007, 0.010]	0.052 [0.043, 0.062]	0.048 [0.040, 0.057]
$\sigma_{\xi, 5}^2$	0.073 [0.062, 0.089]	0.008 [0.007, 0.010]	0.074 [0.061, 0.089]	0.074 [0.060, 0.091]
$\ln \text{MDD}(\mathcal{M}_i   y^T)$	-594.887 (0.166)	-608.410 (0.220)	-575.426 (0.178)	-580.748 (0.131)

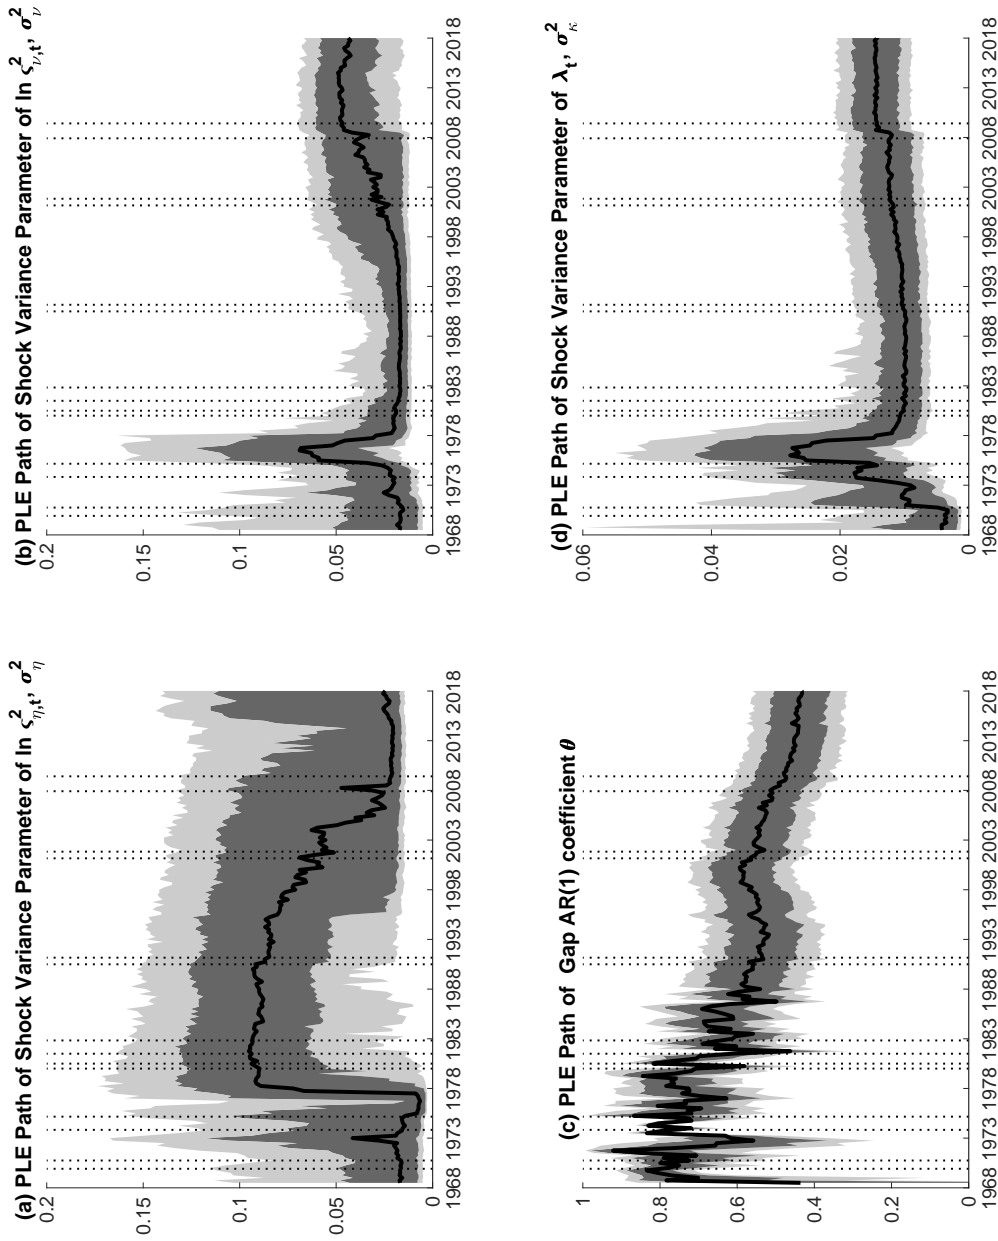
Note: The table contains posterior moments and log MDDs for the state space models  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$  based on  $M = 100,000$  particles and the full data sample. The main entry for every static parameter reports its posterior median with five and 95 percent quantiles in brackets below. Log MDDs for model  $i$  are denoted  $\ln \text{MDD}(\mathcal{M}_i | y^T)$  and computed using equation (17) of the paper. The reported values are the average estimates obtained from 250 repetitions of the particle learning filter, and the associated numerical standard errors appear in parentheses below each estimate.

Figure R.19:  $\mathcal{M}_0$  (w/o noise): Trend and Gap Inflation



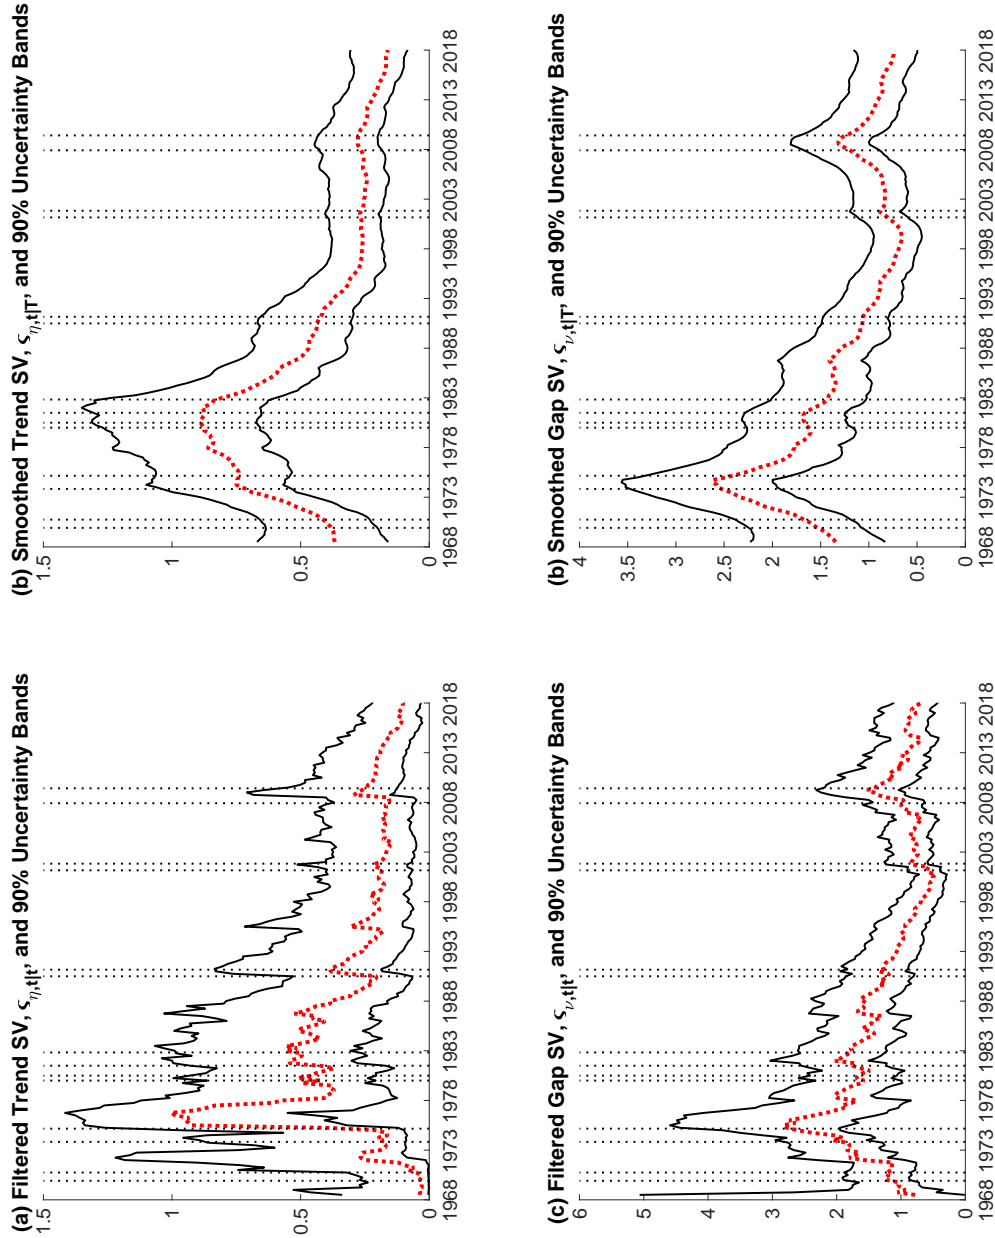
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_0$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.20:  $\mathcal{M}_0$  (w/o noise): Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_0$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.21:  $\mathcal{M}_0$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



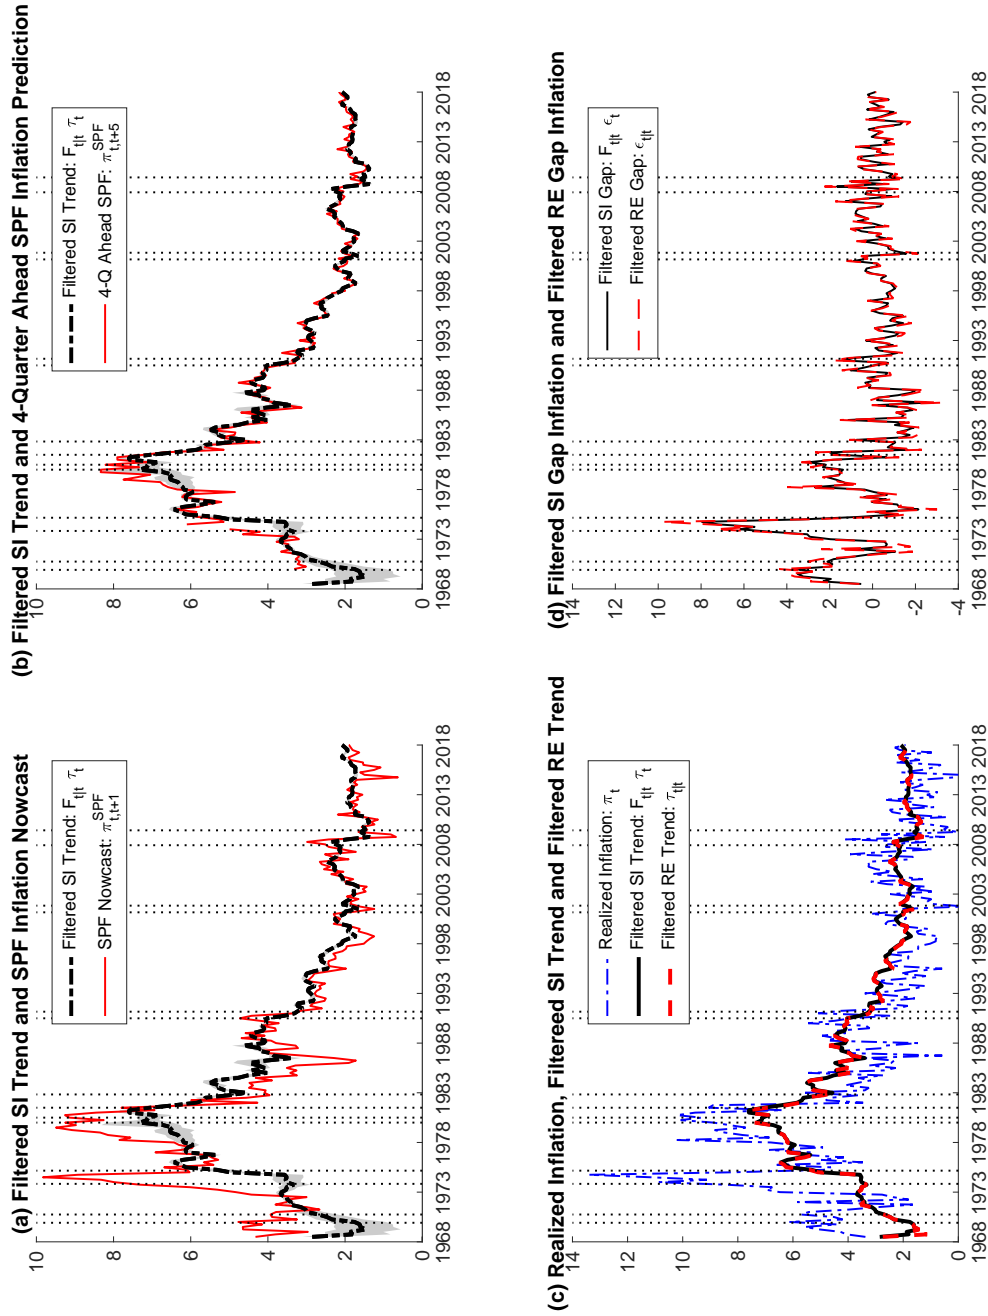
Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_0$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.22:  $\mathcal{M}_0$  (w/o noise): Time-Varying SI Parameter



Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{t|t}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{t|T}$  surrounded by 90% uncertainty bands delineated by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{t|T} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_0$  (w/o noise).

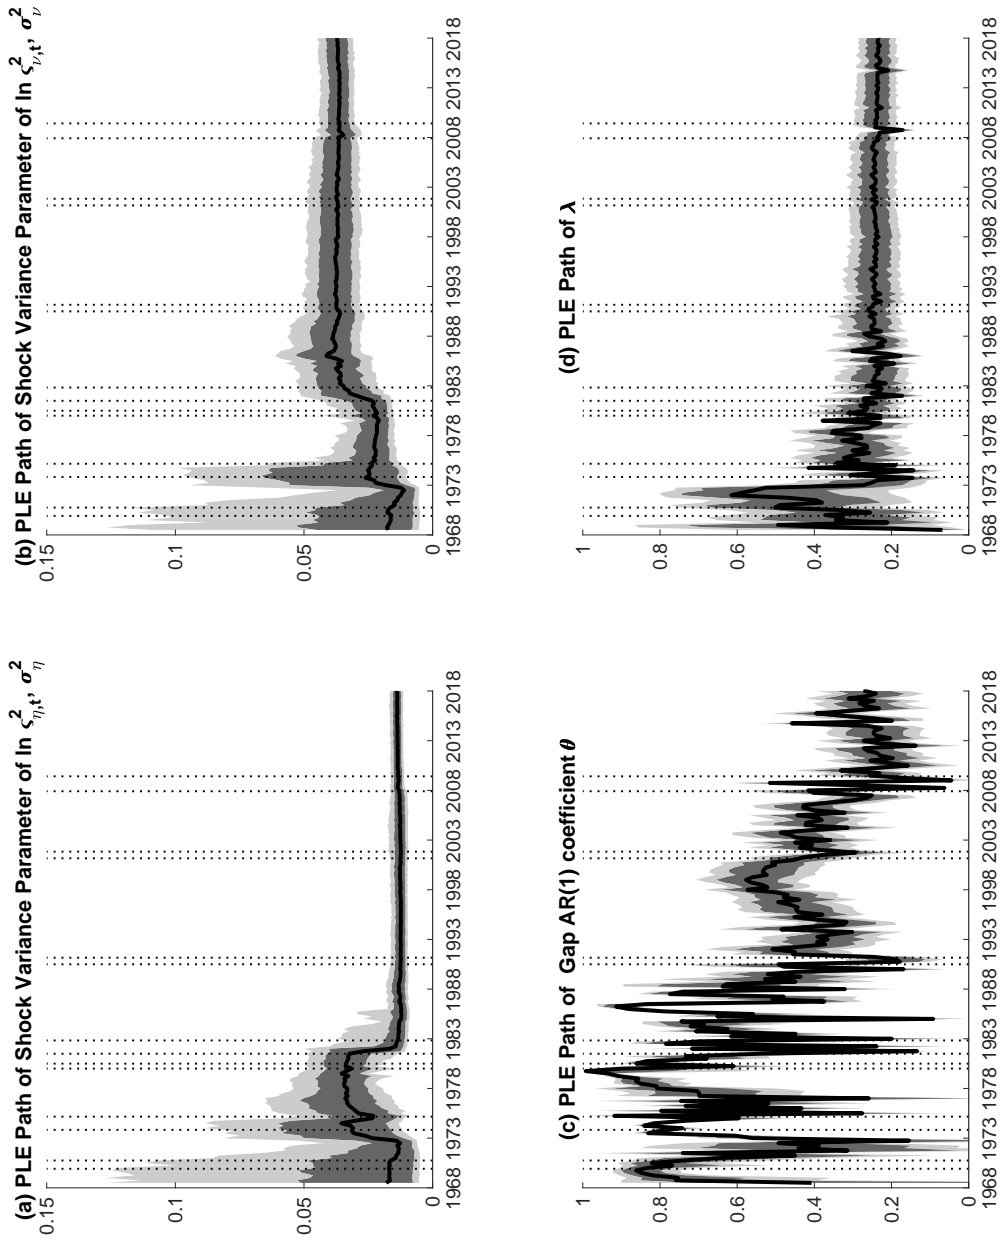
Figure R.23:  $\mathcal{M}_1$  (w/o noise): Trend and Gap Inflation



Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_1$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

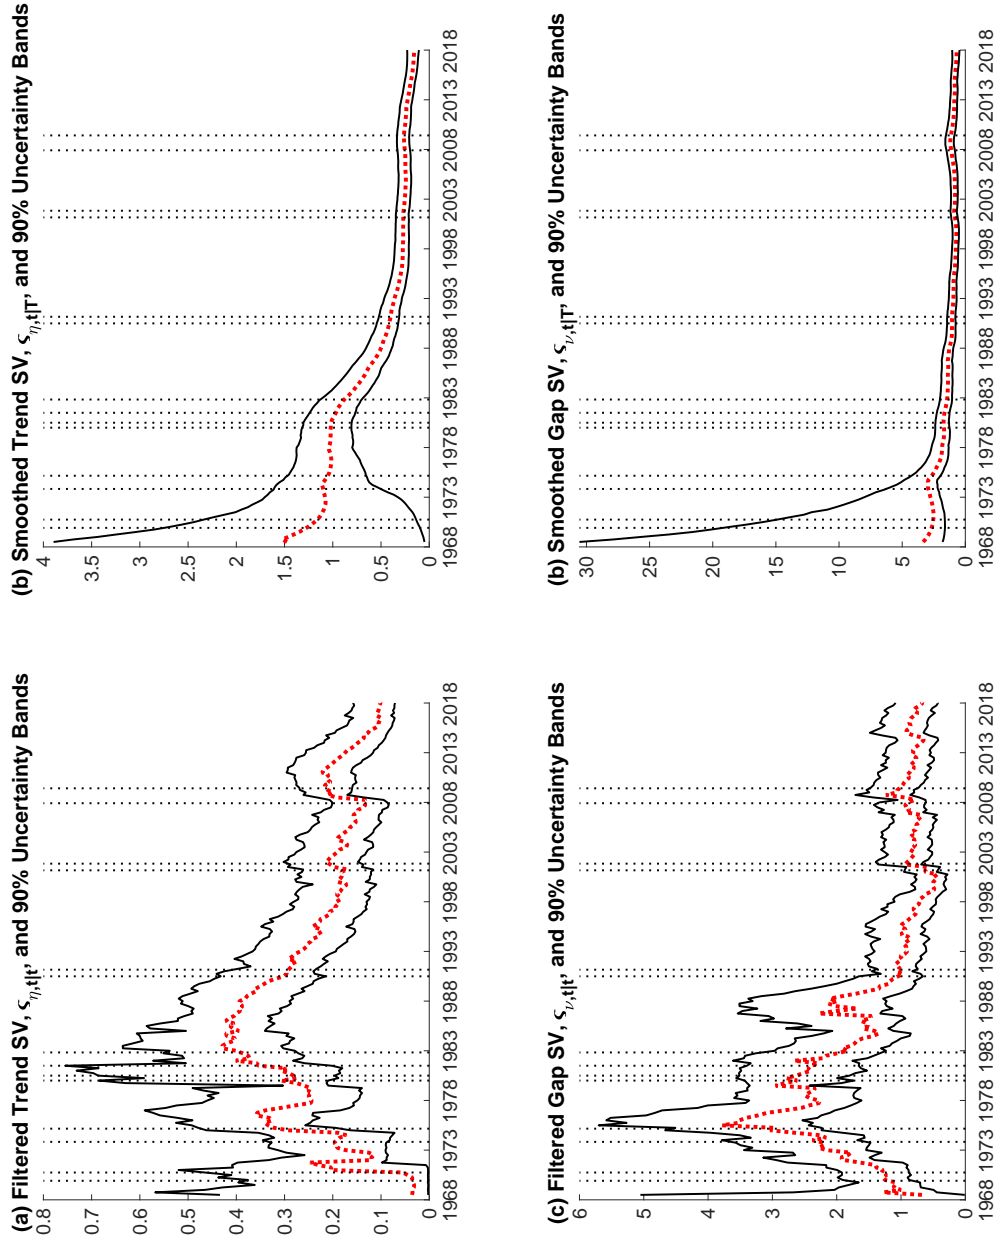


Figure R.24:  $\mathcal{M}_1$  (w/o noise): Static Parameters



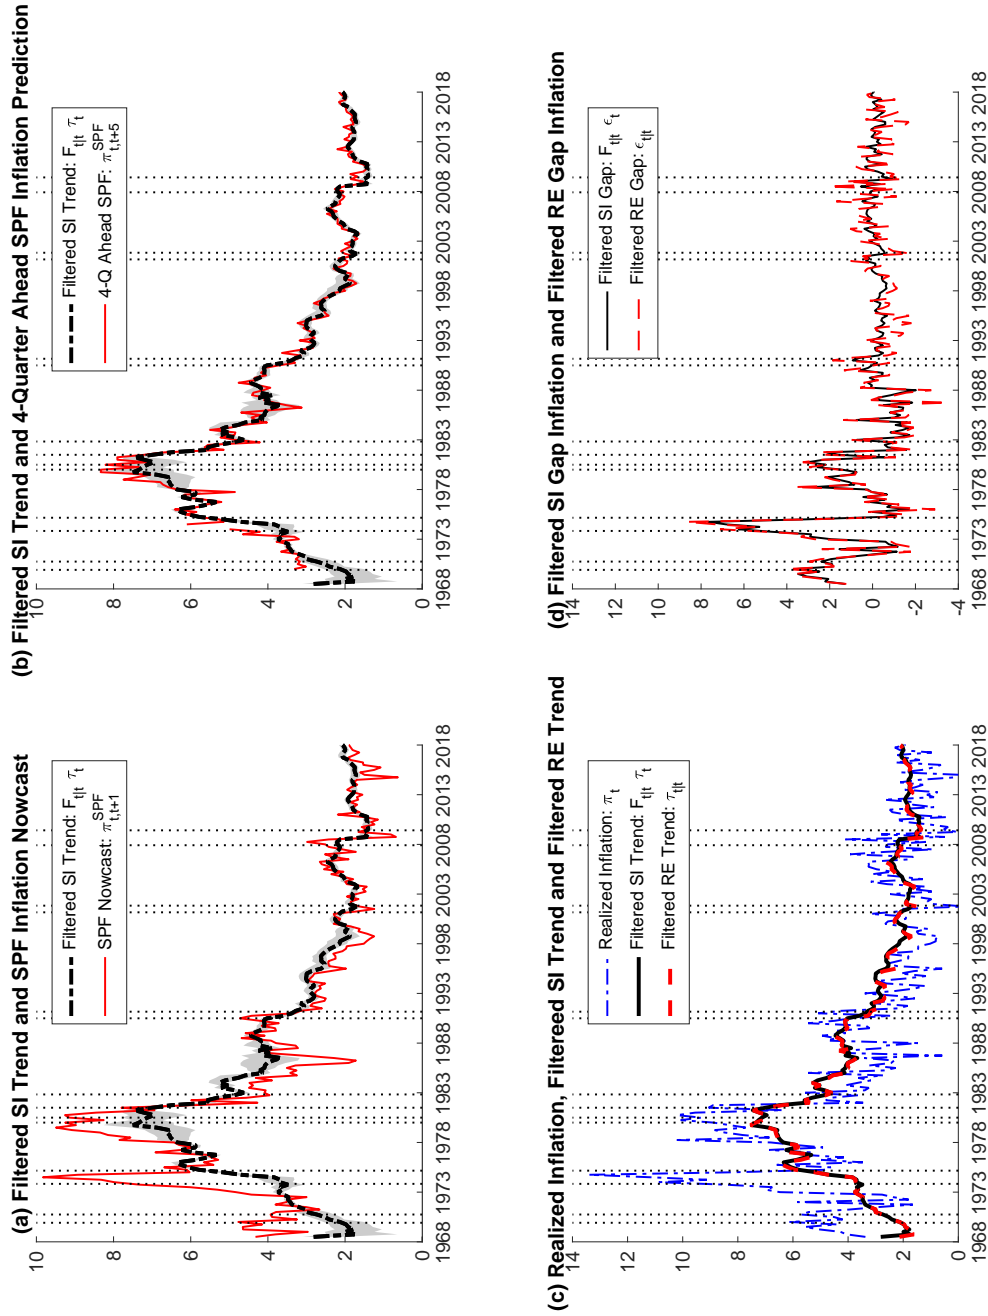
Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_1$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.25:  $\mathcal{M}_1$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



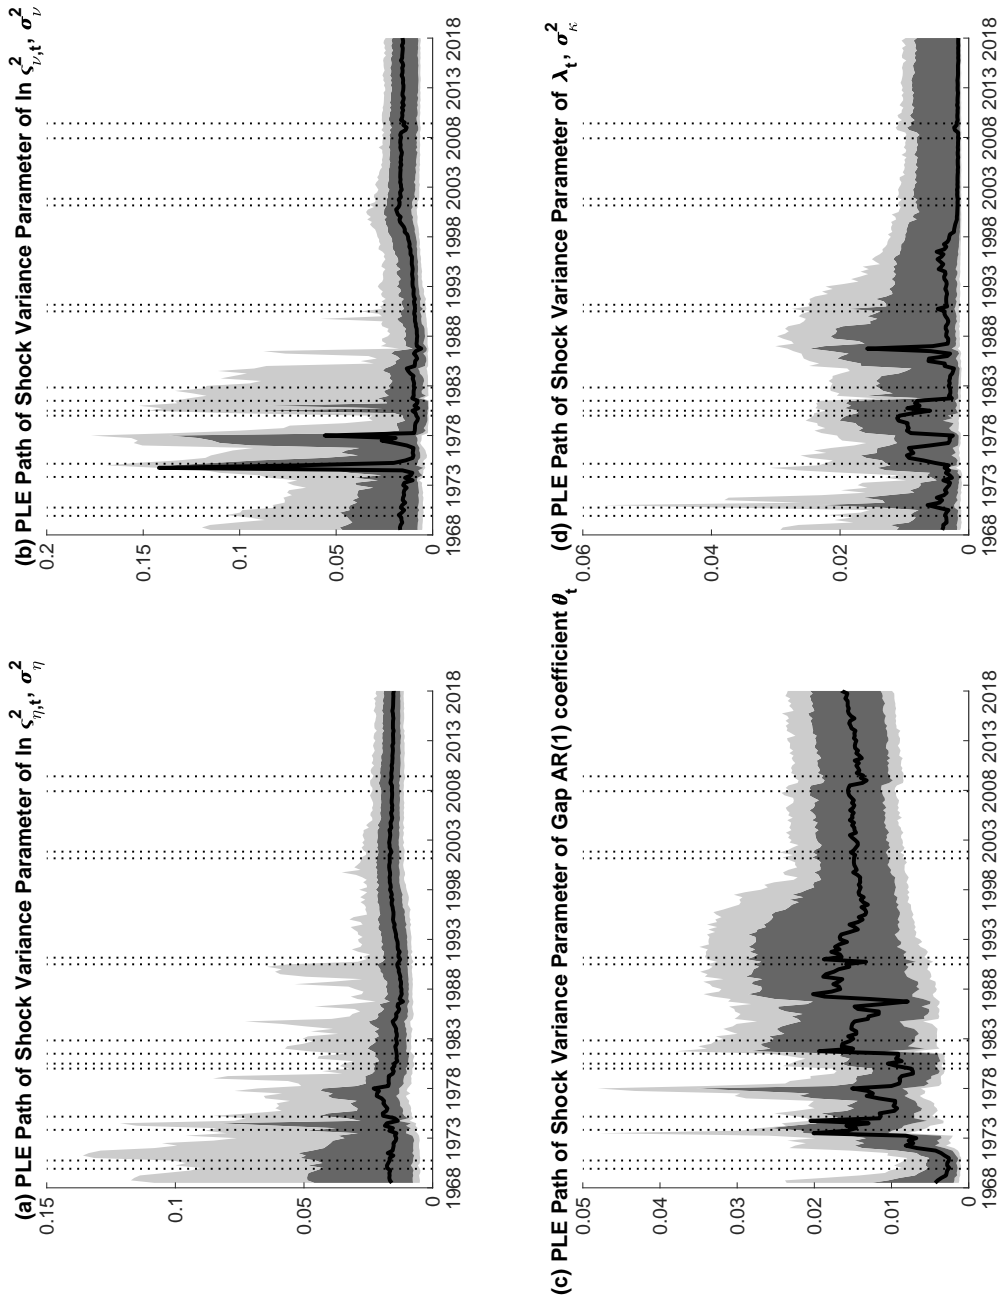
Note: The solid thin (black) lines around estimated and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_1$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.26:  $\mathcal{M}_2$  (w/o noise): Trend and Gap Inflation



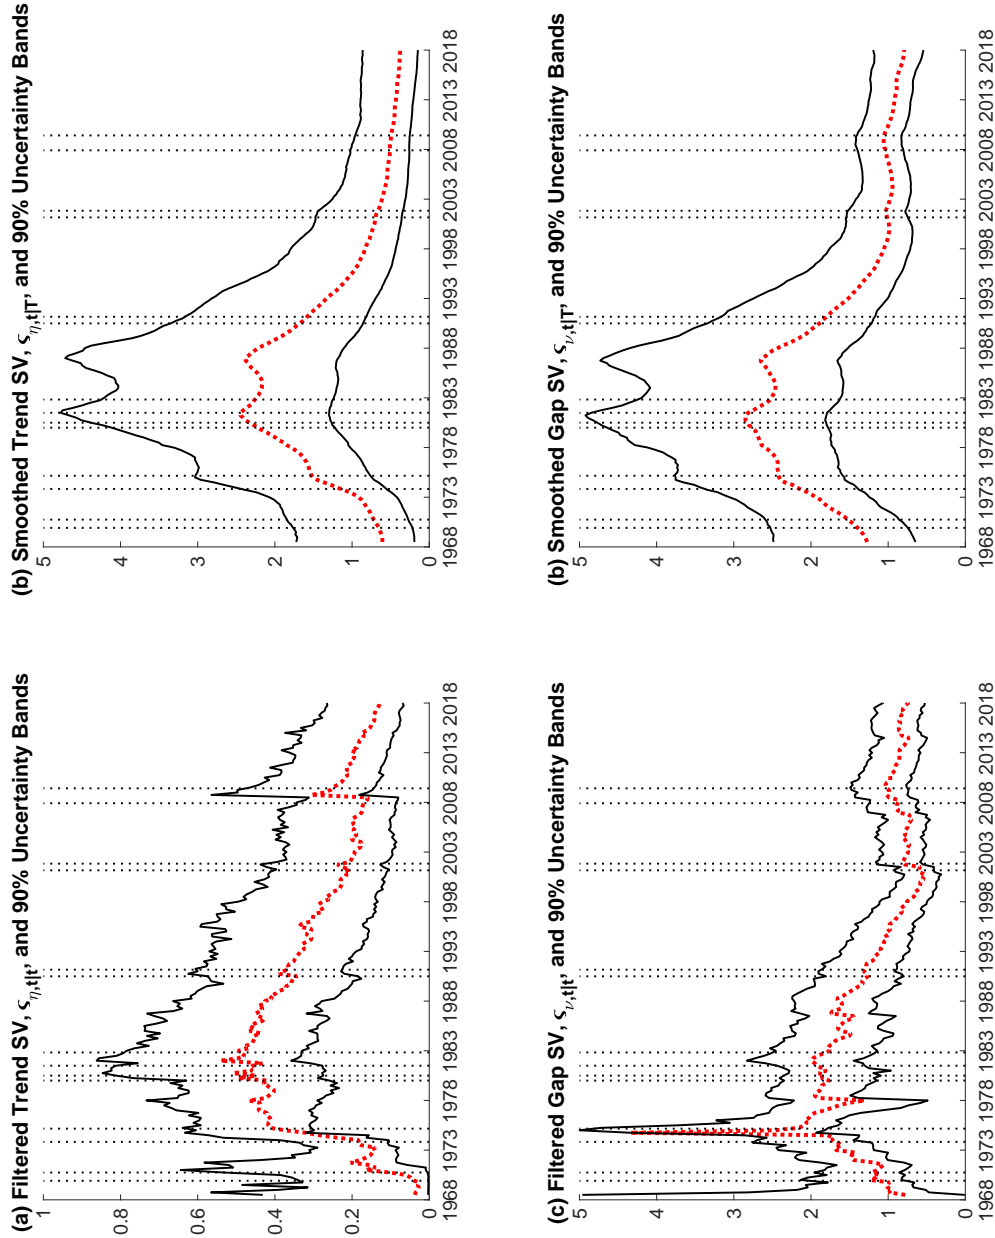
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_2$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.27:  $\mathcal{M}_2$  (w/o noise): Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_2$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.28:  $\mathcal{M}_2$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



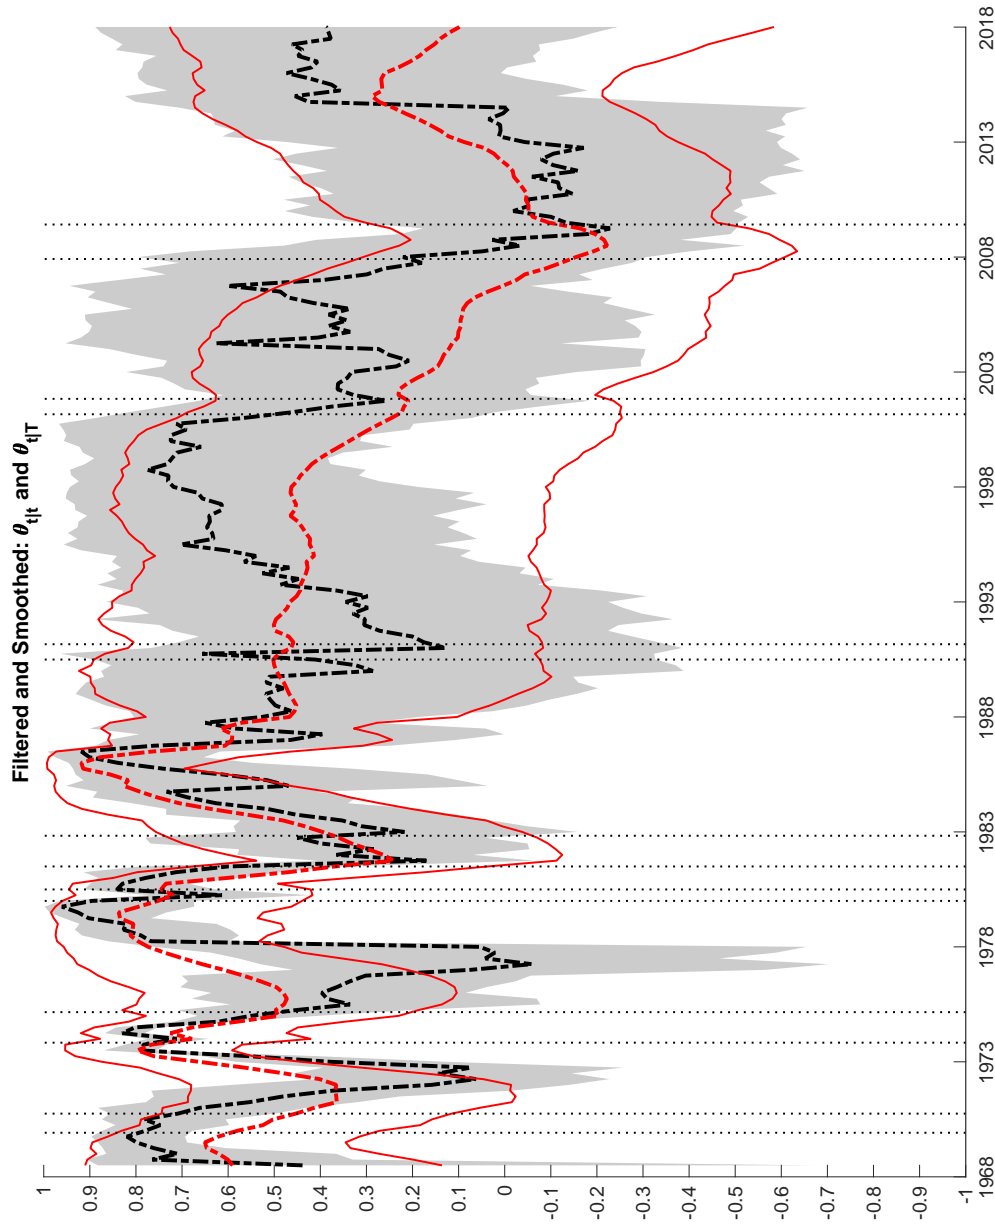
Note: The solid thin (black) lines around estimated and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_2$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.29:  $\mathcal{M}_2$  (w/o noise): Time-Varying SI Parameter



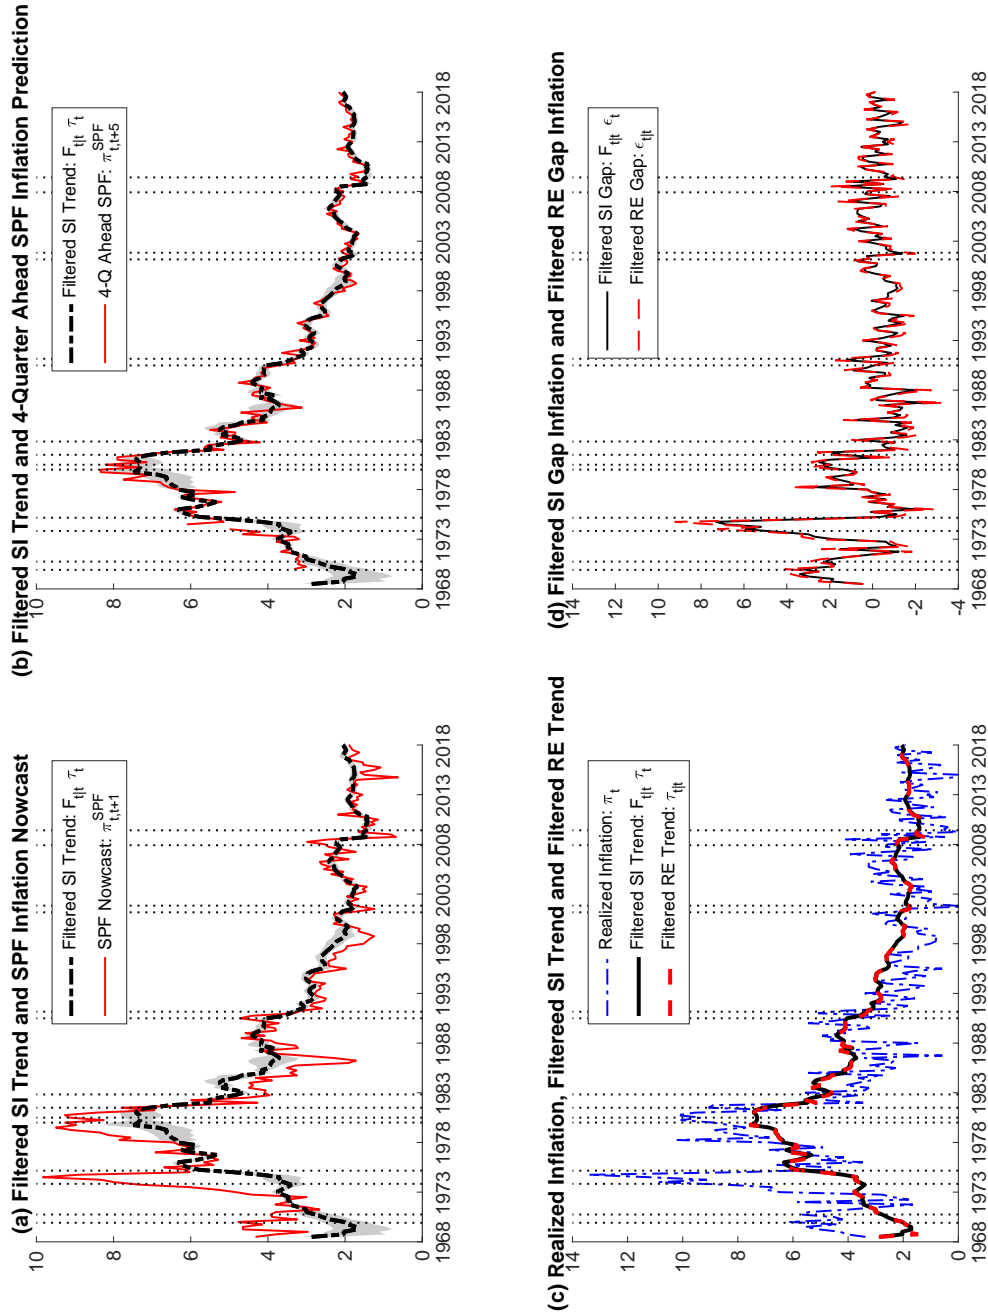
Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{t|t}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{t|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{t|T} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_2$  (w/o noise).

Figure R.30:  $\mathcal{M}_2$  (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process



Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{qt}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{qt}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_2$  (w/o noise).

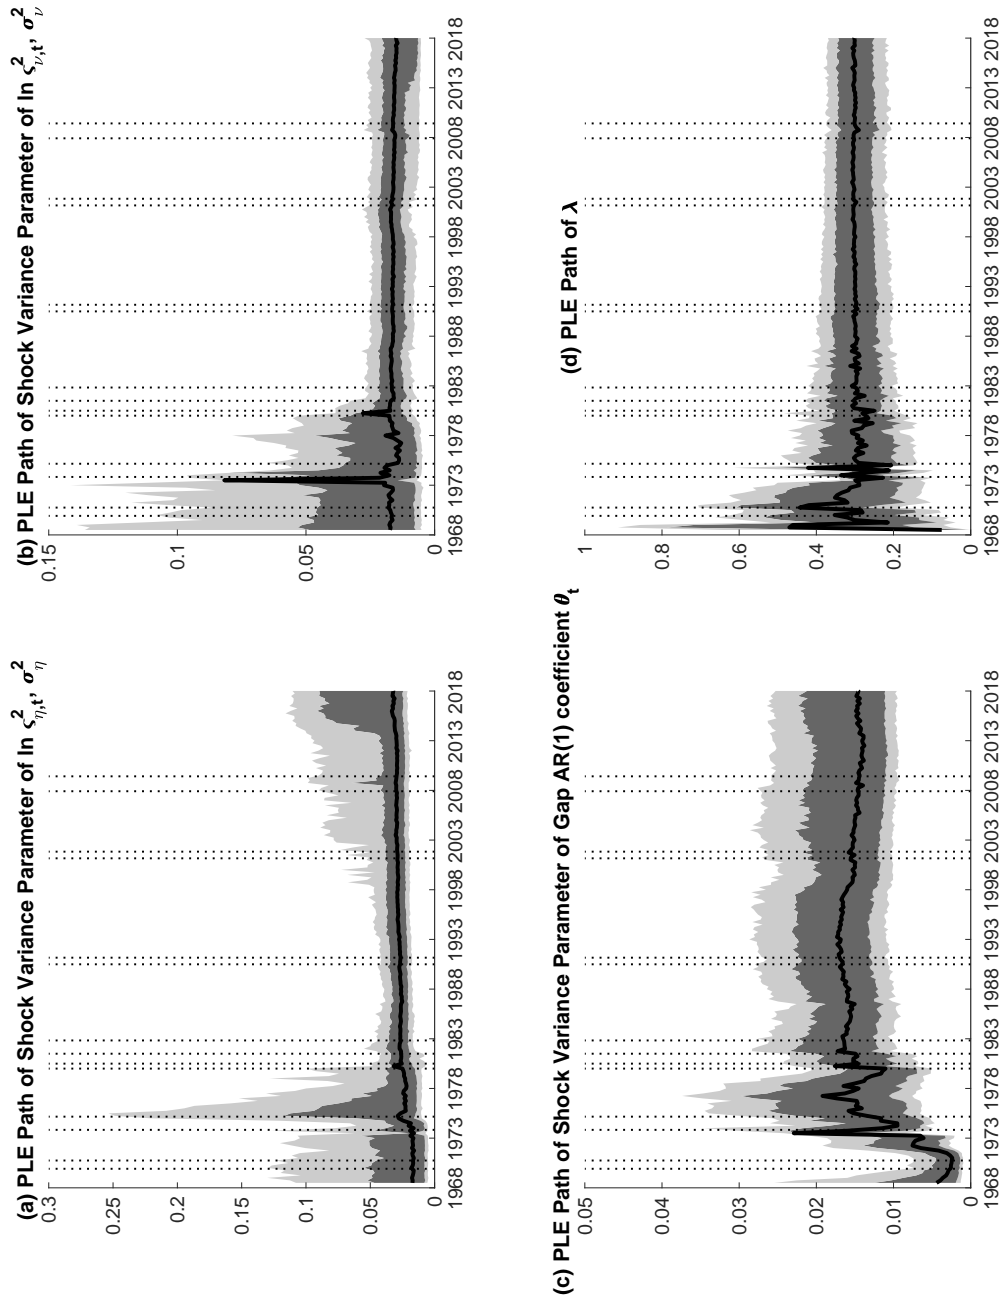
Figure R.31:  $\mathcal{M}_3$  (w/o noise): Trend and Gap Inflation



Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{\text{fit}} \tau_t$  generated by Model  $\mathcal{M}_3$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

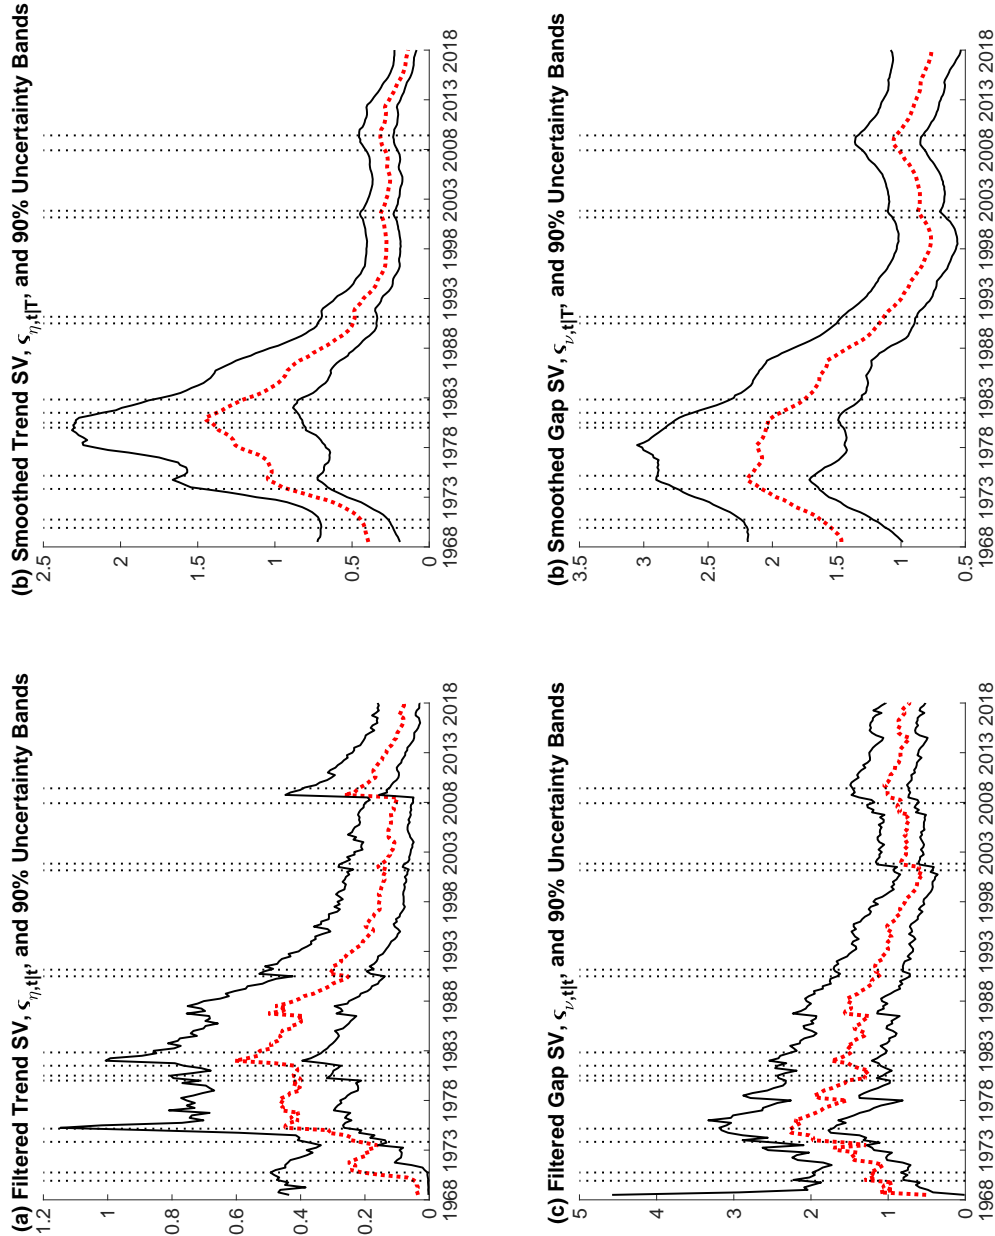


Figure R.32:  $\mathcal{M}_3$  (w/o noise): Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_3$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.33:  $\mathcal{M}_3$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimated and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_3$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.34:  $\mathcal{M}_3$  (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process



Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{f|t}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_3$  (w/o noise).

### R.2.3 Log MDD over time and for different numbers of particles

Figure R.35 displays the evolution of log MDD estimates for our four state space models over time; figure R.36 displays corresponding estimates for the case of no noise in the inflation process. Both figures show that the models where the gap persistence parameter,  $\theta_t$ , is time-varying,  $\mathcal{M}_2$  and  $\mathcal{M}_3$ , accrue substantial contributions to their log MDDs during the Great Inflation period of the 1970s.


Figure R.37 displays the distribution of log MDD estimates generated by  $N = 250$  repetitions of the particle learning filter applied to each of our model variants. For each model, distributions of log MDD estimates are generated for  $M = 100,000$ ,  $M = 10,000$ , and  $M = 1,000$  particles. These distributions represent uncertainty from the Monte Carlo approximation of the true log MDD associated with each model. We gauge the uncertainty of the simulation-based estimates of the log MDD with numerical standard errors. Denoting the log MDD estimate generated by the  $n$ -th simulation  $\mu_n$ , the numerical standard errors around the average estimate,  $\bar{\mu} = \sum_{n=1}^N \mu_n / N$ , are given by

$$\sigma_{\bar{\mu}} = \sqrt{\frac{\sigma_{\mu}^2}{N}} \quad \text{with} \quad \sigma_{\mu}^2 = \frac{1}{N} \sum_{n=1}^N (\mu_n - \bar{\mu})^2 . \quad (\text{R.2})$$

The use of numerical standard errors for gauging the uncertainty of simulation-based estimates is grounded in the work of Geweke (1989), see also Fuentes-Albero and Melosi (2013), and Herbst and Schorfheide (2014) for applications in the context of log MDD estimates.


As reported in Table 5 of the paper, with  $M = 100,000$ , the numerical standard errors are fairly tight, and the mean log MDD estimates are clearly distinguished across models. As shown in Figure R.37, for choices of the number of particles lower than  $M = 100,000$ , the simulated log MDD distributions display more considerable dispersion, fatter tails and even some skew. Figure R.38 presents similar results for the four model variants when noise in the inflation equation is set to zero.

Figure R.35: log MDD over time



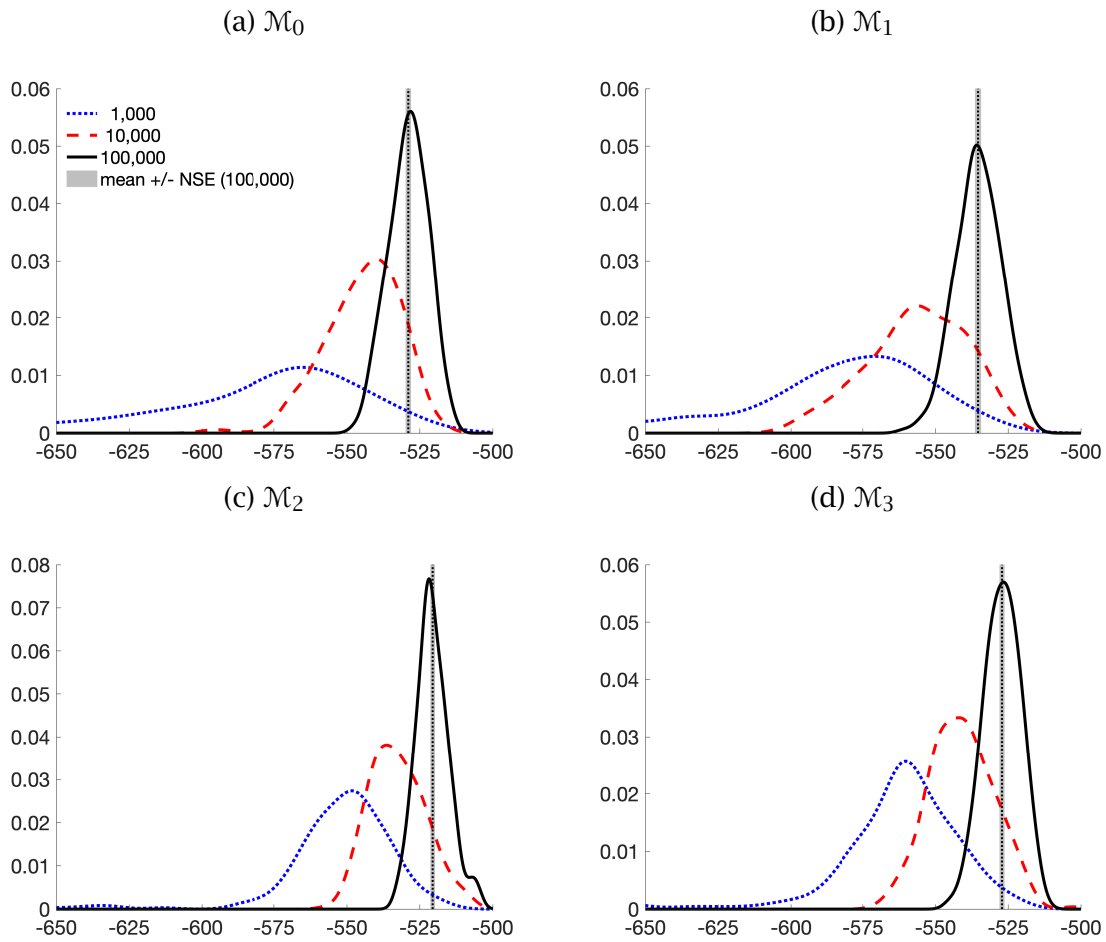
Note: Evolution of log MDD estimates,  $\ln \text{MDD}(\mathcal{M}_i | y^t) \cdot T/t = \sum_{k=1}^t \frac{T}{t} \ln p(y_t | y^{t-1}; \mathcal{M}_i)$ , of each model  $i$  over time. For better comparability, the log MDD are scaled by  $T/t$ . The endpoint of each time series plotted corresponds to the log MDD value tabulated in table 5 of the paper. Log MDDs are computed as average over  $N = 250$  repetitions of the particle learning filter applied to each of our model variants with  $M = 100,000$  particles. The figure omits values for the first eight observations of the sample, which are very low.

Figure R.36: log MDD over time (w/o noise)



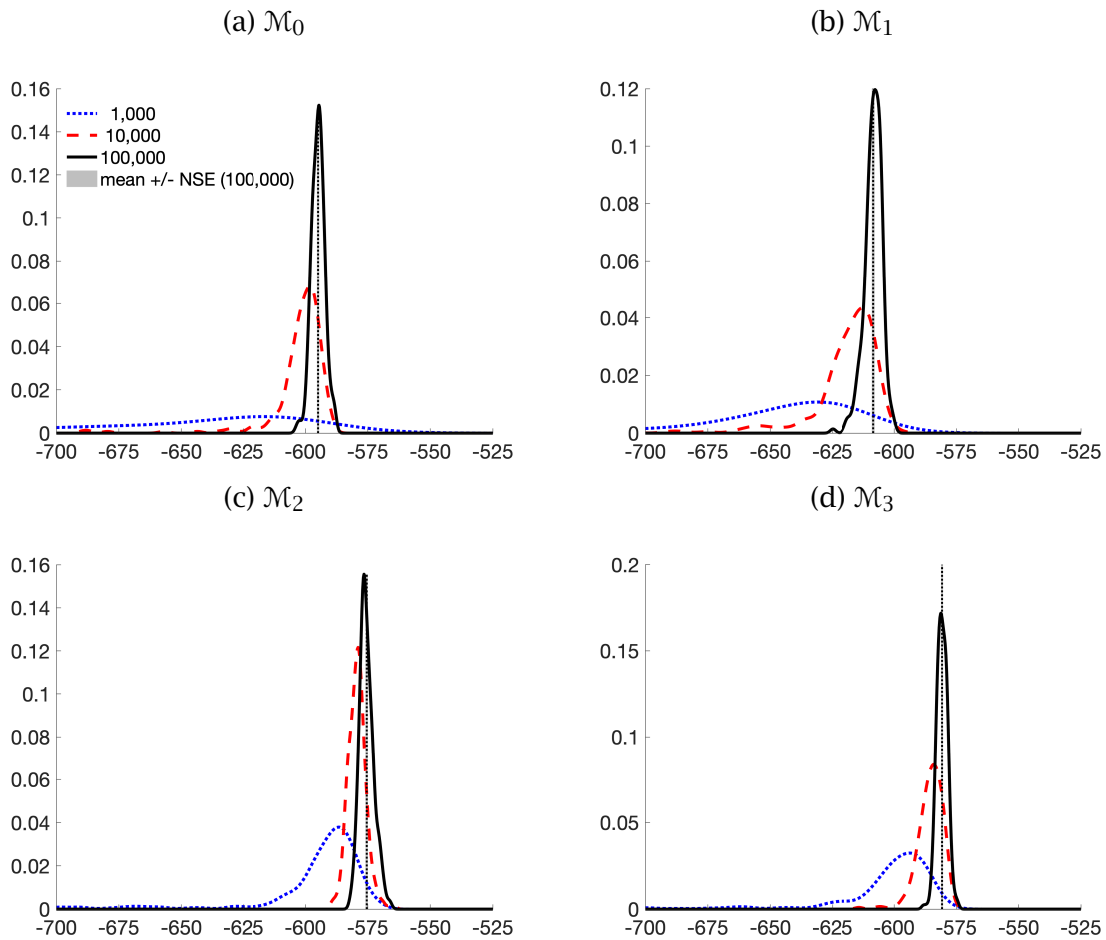
Note: Evolution of log MDD estimates,  $\ln \text{MDD}(\mathcal{M}_i | y^t) \cdot T/t = \sum_{k=1}^t \frac{T}{t} \ln p(y_t | y^{t-1}; \mathcal{M}_i)$ , of each model  $i$  over time. For better comparability, the log MDD are scaled by  $T/t$ . The endpoint of each time series plotted corresponds to the log MDD value tabulated in table R.3 above. Log MDDs are computed as average over  $N = 250$  repetitions of the particle learning filter applied to each of our model variants with  $M = 100,000$  particles. The figure omits values for the first eight observations of the sample, which are very low.

Figure R.37: log MDD for different numbers of particles



Note: Distribution of log MDD estimates generated by 250 repetitions of the particle learning filter applied to each model variant. For each model, distributions of log MDD estimates are generated for  $M = 100,000$  (solid, black),  $M = 10,000$  (dashed, red) and  $M = 1,000$  (dotted, blue) particles. In each case, the distributions shown are kernel density estimates obtained from the 250 simulated log MDD values. The horizontal dashed (black) line, displays the average estimate obtained with  $M = 100,000$  particles, and the surrounding shaded area demarcates bands of plus/minus twice the estimated numerical standard error.

Figure R.38: log MDD for different numbers of particles (w/o noise)



Note: Distribution of log MDD estimates generated by 250 repetitions of the particle learning filter applied to each model variant. For each model, distributions of log MDD estimates are generated for  $M = 100,000$  (solid, black),  $M = 10,000$  (dashed, red) and  $M = 1,000$  (dotted, blue) particles. In each case, the distributions shown are kernel density estimates obtained from the 250 simulated log MDD values. The horizontal dashed (black) line, displays the average estimate obtained with  $M = 100,000$  particles, and the surrounding shaded area demarcates bands of plus/minus twice the estimated numerical standard error.




## R.2.4 Relative effective sample sizes

Figure R.39 reports relative effective sample sizes (ESS) for the particle learning estimates of our four model variants when estimated with 10,000 particles. Relative ESS are computed as  $\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$  where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix.

The ESS are generally quite high, often exceeding 90%. But, during the mid- and late-1970s, and to a lesser degree also during the latest recession, the relative ESS is at times very low, at least for models  $\mathcal{M}_0$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$ , where the gap persistence parameter  $\theta_t$  or the forecast stickiness parameter  $\lambda_t$  are time-varying. Nevertheless, with 10,000 particles, a relative ESS of one percent, i.e. 0.01, still corresponds to an ESS of 1,000 particles. As shown in Figure R.39, the relative ESS does not fall below this threshold, except for a few individual observations during the 1970s for models  $\mathcal{M}_2$  and  $\mathcal{M}_2$ , that are typically associated with particle values for  $\theta_t$  (the AR(1) gap parameter) near the unit circle.

Figure R.40 shows similar results for the relative ESS derived from our four model variants when noise in the inflation equation is set to zero, as in (R.1). (If anything, in the absence of noise in the inflation equation, particularly low values of the relative ESS occur a little more frequently.)

Figure R.39: Relative ESS




Note: Relative effective sample sizes (ESS) with 10,000 particles for the particle learning estimates of our four model variants. Relative ESS are computed as

$$\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$$

where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix. With 10,000 particles, a relative ESS of one percent, i.e. 0.01, corresponds to an ESS of 1,000. Observations of relative ESS below one percent are marked with a diamond (red).

Figure R.40: Relative ESS (w/o noise)



Note: Relative effective sample sizes (ESS) with 10,000 particles for the particle learning estimates of our four model variants when noise in the inflation equation has been shut off, as in (R.1). Relative ESS are computed as

$$\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$$

where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix. With 10,000 particles, a relative ESS of one percent, i.e. 0.01, corresponds to an ESS of 1,000. Observations of relative ESS below one percent are marked with a diamond (red).

### R.3 An alternative specification for time variation in $\lambda_t$

This section presents results obtained from an alternative specification for the process of  $\lambda_t$ , the time-varying SI parameter. The alternative specification adopts an approach used by Del Negro et al (2016) in tracking the time-varying model weight of a dynamic prediction pool. Overall, this alternative specification yields results that are broadly similar to those obtained with our baseline assumption of a bounded random walk for  $\lambda_t$ . In particular, filtered estimates of  $\lambda_t$  display similar time variation as in our original results. However, time variation in smoothed estimates of  $\lambda_t$  is muted compared with estimates found in the paper for  $\mathcal{M}_2$ .

The alternative specification constructs the SI parameter  $\lambda_t \in [0, 1]$  by transforming an unrestricted normally distributed variable, denoted  $x_t^\lambda$ . The transformation employs the cumulative density function of the standard normal distribution, denoted  $\Phi(\cdot)$ , to map  $x_t^\lambda$  into the SI parameter:  $\lambda_t = \Phi(x_t^\lambda)$ . As in Del Negro et al (2016), the dynamics of  $x_t^\lambda$  are described by a stationary AR(1) process:

$$x_t^\lambda = (1 - \rho_x) \mu_x + \rho_x x_{t-1}^\lambda + \sqrt{(1 - \rho_x^2)} \sigma_x \varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad |\rho_x| < 1. \quad (\text{R.3})$$

The process for  $x_t^\lambda$  above has conveniently been parameterized in terms of its unconditional moments. For example, it implies the ergodic distribution for  $x_t^\lambda$  is  $x_t^\lambda \sim N(\mu_x, \sigma_x^2)$ . As discussed by Del Negro et al (2016), the resulting law of motion for  $\lambda_t$  nests a number of interesting special cases: With  $\mu_x = 0$  and  $\sigma_x = 1$ , the ergodic distribution of  $x_t^\lambda$  is standard normal, which after transformation via  $\Phi(\cdot)$  results in a uniform distribution for  $\lambda_t$  over its entire support.<sup>7</sup> With  $\sigma_x > 1$ , the distribution of  $\lambda_t$  becomes U-shaped, indicating that the SI parameter becomes relatively more likely to take values in the tails of its distribution (i.e. values close to zero or one), whereas with  $\sigma_x < 1$ , the distribution of  $\lambda_t$  has an inverted-U shape and is less likely to move away from its mean. Values of  $\mu_x > 0$  shift the mean of  $\lambda_t$  above 0.5 and vice versa. Finally, the AR parameter  $\rho_x$  governs the persistence of  $x_t^\lambda$  and thus also the persistence of  $\lambda_t$ . With  $\rho_x = 1$ ,  $x_t^\lambda$  and  $\lambda$  become constants; with  $\rho_x = 0$   $x_t^\lambda$  and  $\lambda$  are *iid*.

In contrast to Del Negro et al. (2016), who estimate their counterpart to our  $\lambda_t$  with a particle-MCMC procedure, we would like to perform parameter inference with the particle learning filter previously described. To do so, we need to use conjugate prior-posterior relationships to track the posteriors for the parameters by means of sufficient statistics. Inference on AR(1) parameters with conjugate priors is readily available when parameterized in terms of intercept, lag-coefficient and shock variance.<sup>8</sup> However, this is not the case when the AR(1) is parameterized in terms of its unconditional moments as above. Nevertheless, for a given value of  $\rho_x$  it is straightforward to embed inference on  $\mu_x$  and  $\sigma_x^2$  within our particle learning filter using a conjugate normal-inverse-gamma prior as described below. Moreover, using the above parametrization (with a fixed  $\rho_x$ ) lets us directly control the prior for  $\lambda_t$ . Choosing a uniform prior, places equal weight on

<sup>7</sup>For the remainder of this appendix, and unless otherwise noted, we will refer to the ergodic distribution of  $\lambda_t$  simply as “the distribution.”

<sup>8</sup>In this case, the usual normal-inverse-gamma updating for the case of a simple univariate regression would apply; see, for example, Koop (2003).

all permissible values of  $\lambda_t$ , as discussed above. Since we are concerned with persistent changes of  $\lambda_t$ , we have chosen to implement this alternative specification only for a fixed value of  $\rho_x$ , using  $\rho_x = 0.99$ .<sup>9</sup> Our specific calibration of  $\rho_x$  near the unit root reflects our hypothesis that variations in forecast stickiness are slow-moving and long-lasting.<sup>10</sup>

With a given value of  $\rho_x$ , we are left with the problem of embedding inference about the remaining parameters in (R.3), that is  $\mu_x$  and  $\sigma_x^2$ , into the particle learning filter. The goal is to control the priors for both parameters to target a uniform distribution for  $\lambda_t$  over its support between zero and one. Specifically, we use the following Normal-inverse-gamma priors:

$$\sigma_x^2 \sim IG(\alpha_{0,x}, \beta_{0,x}) \quad \text{with } \alpha_{0,x} = 3, \quad \beta_{0,x} = 1, \quad (\text{R.4})$$

$$\mu_x | \sigma_x^2 \sim N(\mu_{0,x}, \sigma_x^2 \cdot V_{0,x}) \quad \text{with } \mu_{0,x} = 0, \quad V_{0,x} = 1. \quad (\text{R.5})$$

This centers the priors for  $\mu_x$  and  $\sigma_x^2$  around values that imply a uniform distribution for  $\lambda_t$ . The normal-inverse-gamma prior is conjugate and can be updated to a normal-inverse-gamma posterior, and,  $\alpha_{t,x}$ ,  $\beta_{t,x}$ ,  $\mu_{t,x}$ , and  $\beta_{t,x}$ , are sufficient statistics that characterize the posterior after  $t$  observations. These sufficient statistics are included in the particle swarm of our particle learning filter.

Given particle draws for  $x_{t-1}^\lambda$  and  $x_t^\lambda$  — and dropping particle superscripts  $i$  for ease of notation — the sufficient statistics are updated as follows: First, we quasi-difference and rescale the AR(1) for  $x_t^\lambda$  in (R.3):

$$y_t^\lambda \equiv \frac{x_t^\lambda - \rho_x x_{t-1}^\lambda}{\sqrt{1 - \rho_x^2}} \quad (\text{R.6})$$

$$= \frac{\mu_x}{\sqrt{1 - \rho_x^2}} + \sigma_x \varepsilon_t \quad (\text{R.7})$$

$$\sim N\left(\frac{\mu_x}{\sqrt{1 - \rho_x^2}}, \sigma_x^2\right), \quad (\text{R.8})$$

$$\text{and let } \mu_{y,t-1} \equiv \frac{\mu_{x,t-1}}{\sqrt{1 - \rho_x^2}}, \quad (\text{R.9})$$

$$V_{y,t-1} \equiv \frac{V_{x,t-1}}{1 - \rho_x^2}. \quad (\text{R.10})$$

The resulting updating equations for the sufficient statistics of the normal-inverse-

<sup>9</sup>We have also experimented with an alternative implementation based on estimating an intercept,  $(1 - \rho_x)\mu_x$ , slope,  $\rho_x$ , and shock variance  $(1 - \rho_x^2)\sigma_x^2$  with this prior. However, we found it difficult to generate meaningful results.

<sup>10</sup>Please recall that for  $\rho_x = 1$ ,  $x_t^\lambda$ , and thus also  $\lambda_t$ , become constant, and the constant- $\lambda_t$  case is studied as part of our set of the four model variants discussed in the paper.

gamma distribution for  $\mu_x$  and  $\sigma_x^2$  after observing time  $t$  data are:<sup>11</sup>

$$V_{y,t} = \frac{1}{V_{y,t-1}^{-1} + 1} \tag{R.11}$$

$$\Rightarrow V_{x,t} = (1 - \rho_x^2) V_{y,t}, \tag{R.12}$$

$$\mu_{y,t} = V_{y,t} (V_{y,t-1}^{-1} \mu_{y,t-1} + y_t^\lambda) \tag{R.13}$$

$$\Rightarrow \mu_{x,t} = \sqrt{1 - \rho_x^2} \mu_{y,t}, \tag{R.14}$$

$$\alpha_{x,t} = \alpha_{x,t-1} + 1, \tag{R.15}$$

$$\beta_{x,t} = \beta_{x,t-1} + (y_t - \mu_{y,t})^2 + (\mu_{y,t} - \mu_{y,t-1})^2 V_{y,t-1}^{-1}. \tag{R.16}$$

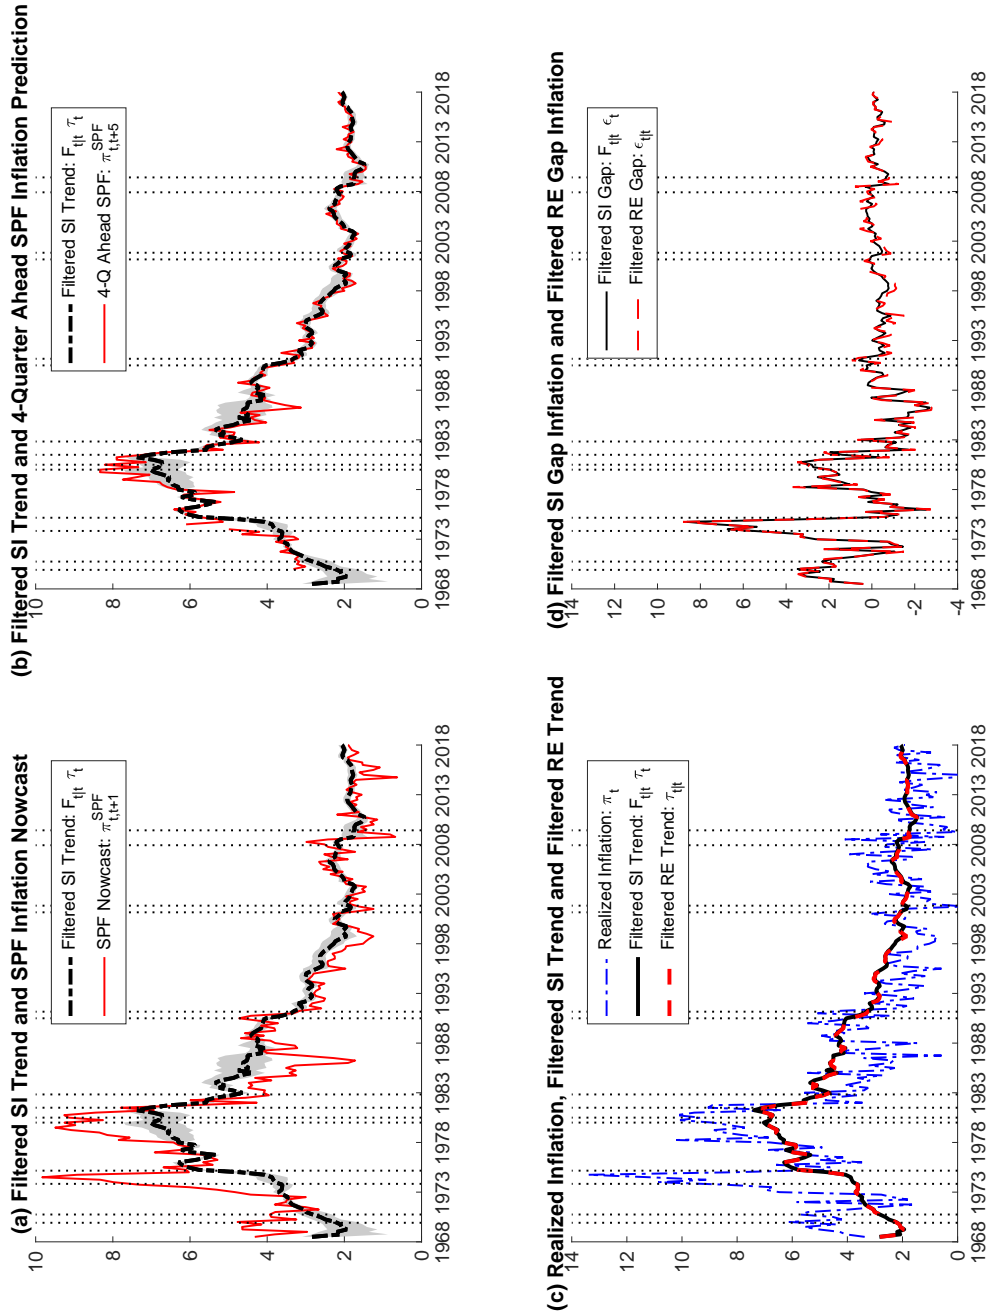
These transition equations for the sufficient statistics of our alternative process of  $\lambda_t$  are straightforward to embed within our particle learning filter described above. Specifically, we estimate a version of model  $\mathcal{M}_2$  using this alternative specification of the process for time-varying  $\lambda_t$ . This alternative version of  $\mathcal{M}_2$  is denoted  $\mathcal{M}_{2,\lambda_x}$ .

As shown in Figures R.41 to R.45, estimates of the latent variables and those parameters that are not specific to  $\lambda_t$ , are broadly similar to what we have found for the estimates of  $\mathcal{M}_2$  discussed in the paper. However, while filtered estimates of  $\lambda_t$  continue to display sizable time variation, smoothed estimates of  $\lambda_t$  display less time variation; in particular, when put in relation to the uncertainty bands surrounding these estimates as shown in Figure R.44.

---

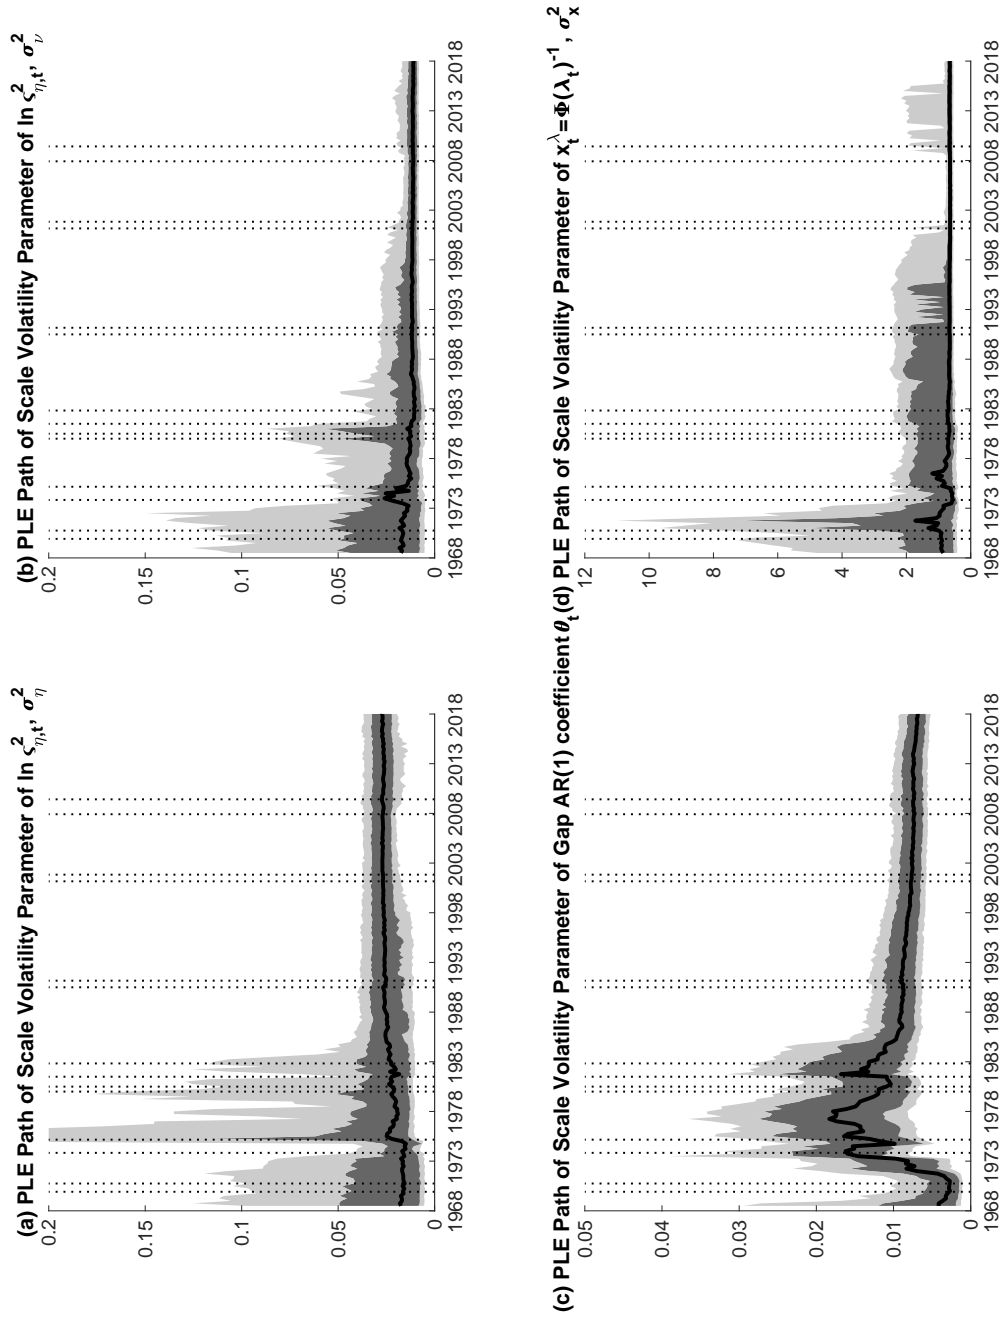
<sup>11</sup>See, for example, the presentation in Koop (2003) (Chapter 3, Exercise 3).

Figure R.41:  $\mathcal{M}_{2,\lambda_x}$ : Trend and Gap Inflation



Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{<t|t} \tau_t$  generated by Model  $\mathcal{M}_{2,\lambda_x}$ . The vertical dotted bands denote NBER dated recessions in the four charts.

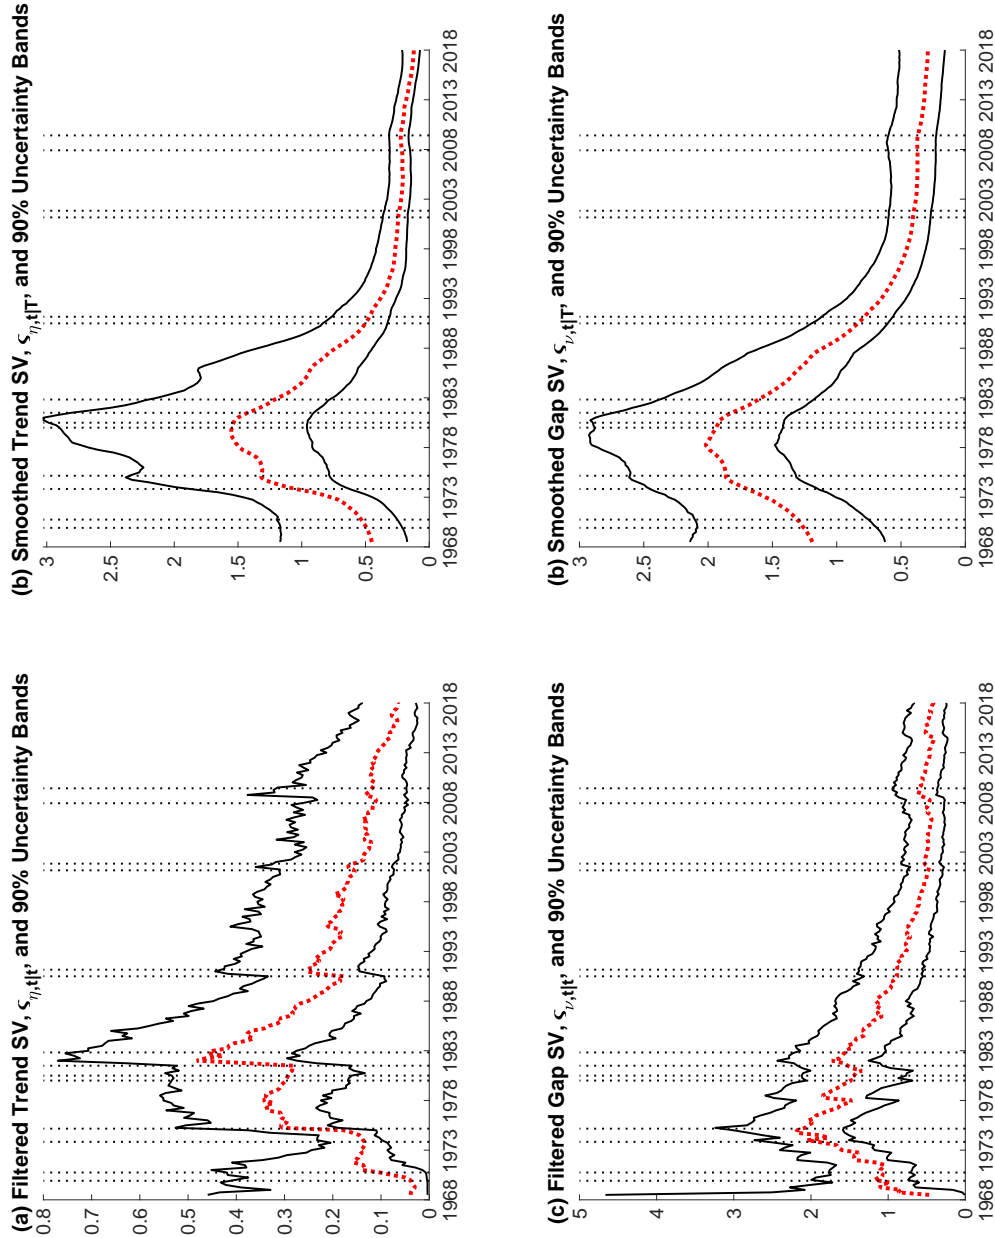
Figure R.42:  $\mathcal{M}_{2,\lambda_x}$ : Static Parameters



Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_{2,\lambda_x}$ . Dotted vertical lines denote NBER recession peaks and troughs.




Figure R.43:  $\mathcal{M}_{2,\lambda_x}$ : Stochastic Volatility in Trend and Gap Inflation




Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_{2,\lambda_x}$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.44:  $\mathcal{M}_{2,\lambda_x}$ : Time-Varying SI Parameter



Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{\text{fit}}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{\text{fIT}}$  surrounded by 90% uncertainty bands delineated by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{\text{fIT}} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_{2,\lambda_x}$ .

Figure R.45:  $\mathcal{M}_{2,\lambda_x}$ : Time-Varying AR Coefficient in the Inflation Gap Process



Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{f|t}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_{2,\lambda_x}$ .

## R.4 An alternative forecast error variance decomposition

In Section 5 of the paper, we employed a decomposition of the forecast error variance into the shares of variations caused by different shocks. Notably, we considered the share of forecast error variances attributable to shocks to trend and persistent gap inflation on the hand — since these shocks affect RE inflation forecasts — and the serially uncorrelated irregular component in inflation on the other hand.<sup>12</sup> As documented in the paper, the share of forecast error variances due to noise shocks has markedly risen since the onset of the Volcker disinflation, broadly mirroring the increase in estimated stickiness as measured by  $\lambda_t$ .

In this section, we briefly describe an alternative measure of inflation persistence, related to variance ratios that Cogley et al (2010) used to characterize persistence in gap inflation. Adapted to our model, these measures indicate a decline in inflation persistence that is similar to what is suggested by the variance share due to noise shocks discussed in the paper. In their application, Cogley et al measure persistence in gap inflation at a particular forecast horizon as the share of the gap variance due to variations in conditional expectations for that horizon. In the context of our application, we are interested in characterizing persistence in overall inflation, not just gap inflation, since SPF forecasts are formulated for overall inflation. Given the presence of the non-stationary trend, there is no well-defined unconditional variance of inflation. Instead, we decompose the forecast error variance of inflation for a given horizon  $h$  into the share of variations attributable to updated conditional expectations next period and the then remaining conditional variance. Applying the law of total variance, we get:

$$\mathbf{Var}_t (\pi_{t+1+h}) = \mathbf{Var}_t (\mathbf{E}_{t+1} \pi_{t+1+h}) + \mathbf{E}_t \mathbf{Var}_{t+1} (\pi_{t+1+h}). \quad (\text{R.17})$$

An alternative measure of inflation persistence is given by

$$R_h^2 = \frac{\mathbf{Var}_t (\mathbf{E}_{t+1} \pi_{t+1+h})}{\mathbf{Var}_t (\pi_{t+1+h})} = 1 - \frac{\mathbf{E}_t \mathbf{Var}_{t+1} (\pi_{t+1+h})}{\mathbf{Var}_t (\pi_{t+1+h})}. \quad (\text{R.18})$$


Figure R.46 provides estimates of this alternative persistence measure for various forecast horizons. Unsurprisingly (and similar to Cogley et al), variations in conditional expectations account for a smaller share of forecast error variance for longer forecast horizons. Nevertheless, across all horizons, a familiar pattern emerges, documenting a clear decline in inflation persistence with the onset of the Volcker inflation (and a very mild, transitory uptick around the Great Recession), which broadly mirrors the contours of our estimates for the evolution of forecast stickiness  $\lambda_t$ .

---

<sup>12</sup>Please recall that the inflation process, as introduced in equations (1) and (2) of the paper, is the sum of three mutually uncorrelated components (trend, persistent gap, and the irregular gap component):  $\pi_t = \tau_t + \varepsilon_t + \sigma_{\zeta, \pi} \zeta_{\pi, t}$  with  $\zeta_{\pi, t} \sim N(0, 1)$ . Passing the conditional variance operator through yields  $\mathbf{Var}_t (\pi_{t+h}) = \mathbf{Var}_t (\tau_{t+h}) + \mathbf{Var}_t (\varepsilon_{t+h}) + \sigma_{\zeta, \pi}$ .

Figure R.46: Alternative  $R^2$  measure of inflation persistence

(a)  $h = 1$



(b)  $h = 4$



Note: Estimates of  $R^2_h$ , as defined in equation (R.18), generated from smoothed posterior of the  $\mathcal{M}_2$  model. Thick (thin) dashed lines denote 68% (90%) uncertainty bands.

## R.5 SPF predictions and trend inflation uncertainty

Figure R.47 compares uncertainty around trend inflation estimates generated from different sets of conditioning variables: inflation alone, a longer-horizon SPF forecast ( $h = 5$ ), the combination of both, or our original data vector consisting of inflation and SPF forecasts for  $h = 1, 2, \dots, 5$ . Uncertainty is measured by the volatility of the filtered posterior for trend inflation, that is  $\sqrt{\text{Var}(\tau_t | \mathcal{Y}^t, \hat{\mathcal{V}}_{t|T}, \hat{\Psi})}$  with  $\hat{\Psi}$  and  $\hat{\mathcal{V}}_{t|T}$  fixed at levels that are described next. For better comparability, while varying the elements of  $\mathcal{Y}^t$  for this exercise, we fix parameters and non-linear state variables at their posterior medians obtained from a full-sample estimation of our preferred model  $\mathcal{M}_2$ . These estimates are denoted  $\hat{\Psi}$  and  $\hat{\mathcal{V}}_{t|T}$ . While conditional on  $\hat{\Psi}$  and  $\hat{\mathcal{V}}_{t|T}$ , the posterior underlying our uncertainty measure is filtered, insofar its inference about trend inflation reflects only data  $\mathcal{Y}^t$  up to  $t$ , not  $T$ .<sup>13</sup>

Figure R.47 plots the posterior volatilities of  $\tau_t$ .<sup>14</sup> In the figures, the solid (black), dashed (gray), dotted (blue), and dot-dashed (red) lines are the volatility of trend inflation conditional on different specifications of  $\mathcal{Y}^t$ : our original measurement vector described above, only realized inflation, only the 5-quarter ahead SPF inflation forecast, and the combination of realized inflation and the 5-quarter ahead SPF forecast.


The key message from Figure R.47 is that  $\pi_{t,t+5}^{SPF}$  contributes the bulk of the information pertinent for estimating  $\tau_t$ . The uncertainty around trend estimates obtained from  $\pi_{t,t+5}^{SPF}$  (dot-dashed line) is mostly not much wider than when estimates are conditioned on the full measurement vector (solid). Nevertheless, conditioning merely on  $\pi_{t,t+5}^{SPF}$  generates somewhat higher uncertainty than the full information set during the double-dip recessions of the early 1980s that is manifested as a hump in the dot-dashed line of figure R.47. This additional uncertainty can, however, be largely reduced by merely adding realized inflation to  $\pi_{t,t+5}^{SPF}$  in the conditioning set (the dashed line).

In contrast, absent any survey forecasts, there is substantially wider uncertainty in trend estimates that condition only on realized inflation (dotted line), at least from 1968 through the 1980s. Prior to the Volcker disinflation there is insufficient information in  $\pi_t$  alone to estimate  $\tau_t$  efficiently. Echoing the discussion of Faust and Wright (2013) about the merits of centering inflation forecasts around long-horizon survey predictions, we conclude  $\pi_{t,t+5}^{SPF}$  has useful information for lowering uncertainty surrounding estimates of  $\tau_t$ .

<sup>13</sup>Conditional on parameters and trajectories for the latent variables, the variance computations follow directly from standard Kalman filtering formulas, as described, for example, in our supplementary appendix.

<sup>14</sup>As shown in earlier working paper versions of this paper, we obtain very similar uncertainty measures for  $\mathbf{F}_t \tau_t$ .

Figure R.47: Uncertainty around trend inflation estimates



Note: Estimates of  $\sqrt{\text{Var}(\tau_t | y^t, \hat{v}_{t|T}, \hat{\Psi})}$  for different choices of  $y_t$ , as indicated in the legend above. Estimates are generated from  $\mathcal{M}_2$  using data from 1968Q4 to 2018Q3.  $\hat{\Psi}$  and  $\hat{v}_{t|T}$  denote full-sample estimates of parameters and non-linear states obtained from our full data vector, comprising inflation and SPF forecasts for  $h = 1, 2, \dots, 5$ . Vertical dotted bands denote NBER dated recessions.

## R.6 Univariate estimates from the Stock-Watson UCSV model

This section of the appendix presents estimates of the unobserved components model with stochastic volatility (UCSV) by Stock and Watson (2007) when applied to our inflation data.<sup>15</sup> The UCSV model is nested as special case within our state space models. In the UCSV model, the inflation gap is serially uncorrelated ( $\theta = 0$ ) and there is no further noise component the measurement equation of inflation. In short, the model is given by:

$$\pi_t = \tau_t + \varepsilon_t \tag{R.19}$$

$$\tau_{t+1} = \tau_t + \varsigma_{\eta,t+1}\eta_t, \quad \eta_t \sim N(0, 1) \tag{R.20}$$

$$\varepsilon_{t+1} = \varsigma_{v,t+1}\nu_t, \quad \nu_t \sim N(0, 1) \tag{R.21}$$

$$\ln \varsigma_{\ell,t+1}^2 = \ln \varsigma_{\ell,t}^2 + \sigma_{\ell}\xi_{\ell,t+1}, \quad \xi_{\ell,t+1} \sim N(0, 1) \quad \ell = \eta, \nu. \tag{R.22}$$

There is only a single measurement equation, R.19, and estimates of the Stock-Watson model are univariate in the sense that they are conditioned only on inflation data. Among others, variants of the UCSV model are estimated by Grassi and Proietti (2010), Stock and Watson (2010), Creal (2012), Shephard (2013), and Cogley and Sargent (2015).

We estimate the model using our particle learning filter, which also provides inferences on the variance parameters of the SV processes. We use the same priors as for the corresponding parameters and initial conditions as in the state space models described in the paper. The particle learning filter has been estimated with  $M = 10,000$  particles.

Figure R.48 and R.49 present estimates of inflation trend and gap and the associated stochastic volatilities. The contours of these estimates are familiar from the prior literature as cited above. The univariate trend estimates track realized inflation quite closely and the surrounding uncertainty are quite a bit wider than those generated by our multivariate state space models. Figure R.50 displays the particle learning estimates of the variance parameters in the SV processes. Finally, figure R.51 shows the relative effective sample sizes of the particle learning filter, which is nearly perfect.


---

<sup>15</sup>As discussed in section 4 of the paper, our inflation measure is given by second-release readings of log differences in the GNP/GDP deflator as provided by the Real-Time Data Set for Macroeconomists of the Federal Reserve Bank of Philadelphia.




Figure R.48: Stock-Watson UCSV model: Inflation Trend and Gap

(a) Inflation Trend




(b) Inflation Gap




Note: Filtered estimates of inflation trend and gap in the Stock-Watson (2007) UCSV model. The estimates are obtained from a particle learning filter. Shaded areas in panel (a) depict 90% uncertainty bands around the trend estimates.

Figure R.49: Stock-Watson UCSV model: SV in Inflation Trend and Gap

(a) Inflation Trend SV




(b) Inflation Gap SV




Note: Filtered estimates of stochastic volatility in shocks to inflation trend and gap in the Stock-Watson (2007) UCSV model. The estimates are obtained from a particle learning filter. Thin lines depict 90% uncertainty bands.

Figure R.50: Stock-Watson UCSV model: Parameter estimates

(a) Variance of shocks to Inflation Trend SV




(b) Variance of shocks to Inflation Gap SV



Note: Paths of particle learning estimates of shock variance parameters in the Stock-Watson (2007) UCSV model. The estimates are obtained from a particle learning filter. Thin lines depict 90% uncertainty bands.

Figure R.51: Stock-Watson UCSV model: Effective Sample Size



Note: Relative ESS of the particle learning filter applied to the Stock-Watson (2007) UCSV model with  $M = 10,000$  particles. Relative ESS are computed as

$$\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$$

where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix.

## R.7 Results with CPI data

This section complements the estimates presented in the paper (and appendix R.2) for inflation measures based on the GNP/GDP deflator with corresponding results based on CPI data. As before, forecast data has been obtained from the SPF. However, the available SPF data extends back only to 1981 and our sample runs from 1981Q2 to 2018Q3. CPI is generally not revised (apart from updates in seasonal adjustments) and our measure of realized CPI inflation was taken from the current vintage of FRED (the St. Louis Federal Reserve Economic Data) per February 19 2019.

Section R.2.1 shows results for our four model variants  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$  where, in each model, the inflation process contains a serially-uncorrelated irregular component, as in equations (1) and (2) of the paper. Section R.2.2 then turns to results based on an alternative specification for inflation that omits the irregular component. Sections R.7.3, and R.7.4 report log MDDs for different numbers of particles and effective sample sizes. Section R.7.5 presents forecast comparison tables.

### R.7.1 Results for the four model variants

Table R.4 reports estimates of static parameters as well as logMDDs for each of the four model variants when noise in the inflation equation has been removed. In comparison to results from our baseline data where inflation measures and forecasts reflect the GNP/GDP deflator, models  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ , and  $\mathcal{M}_2$ , fit the data almost equally well in terms of log MDDs, with  $\mathcal{M}_3$  not too far behind. We attribute this similarity at least in part to the substantially shorter data sample, which largely omits the Great Inflation years.

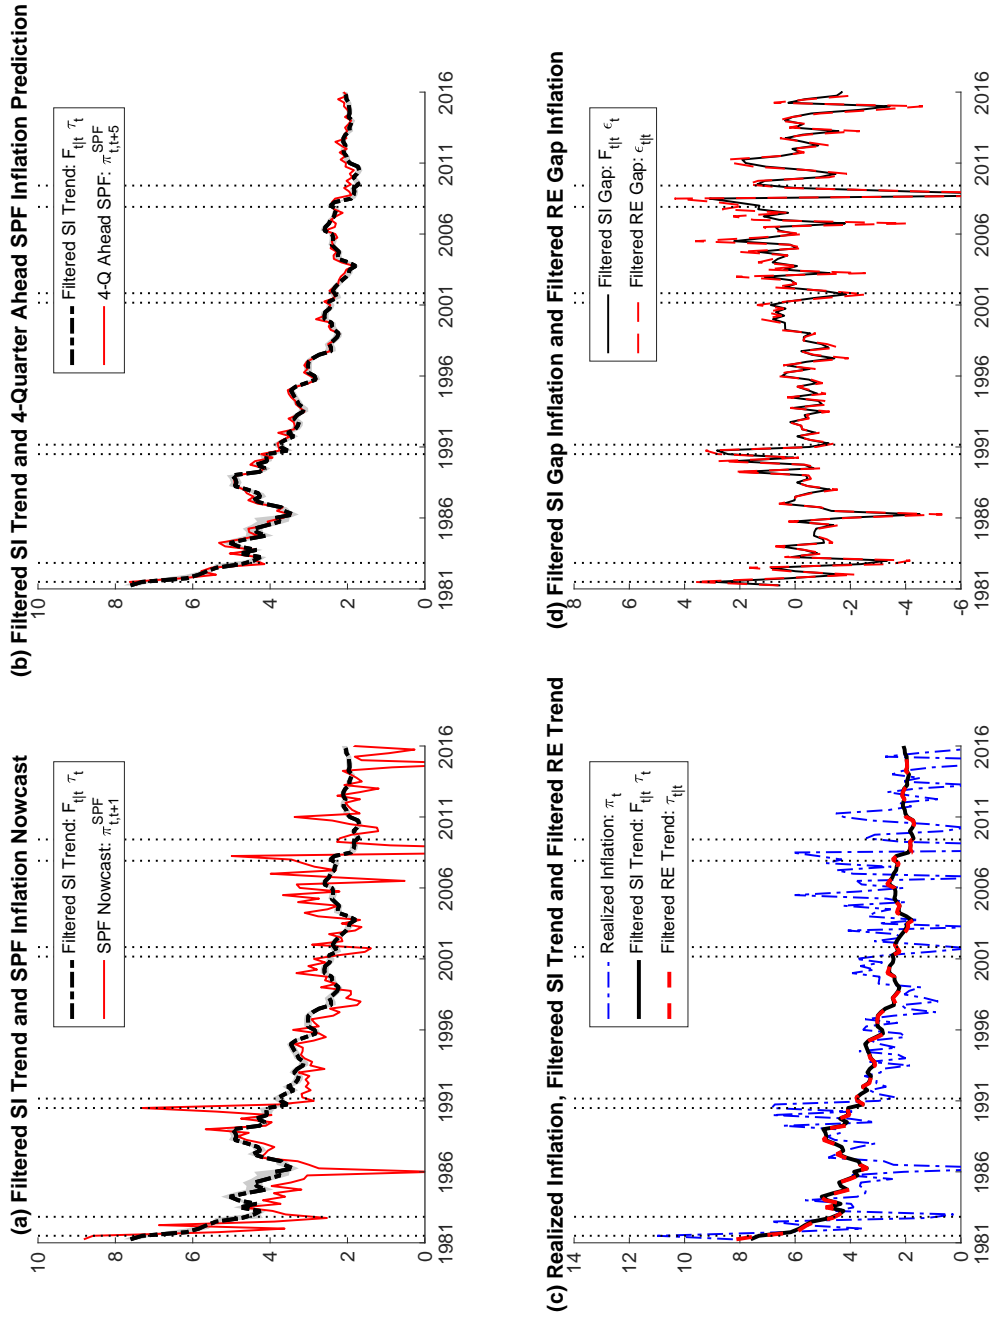
Figures R.3-R.6 report results for  $\mathcal{M}_0$ , with time-varying  $\lambda_t$  but constant  $\theta$ . Figures R.56-R.9 show results for  $\mathcal{M}_1$ , where both  $\lambda$  and  $\theta$  are assumed to be constant. Figures R.10-R.14 provide results for  $\mathcal{M}_2$ , where both  $\lambda_t$  and  $\theta_t$  are time-varying. Finally, Figures R.15-R.18 show results for  $\mathcal{M}_3$ , where  $\theta_t$  is time-varying but  $\lambda$  is assumed constant.

Table R.4: Parameter Estimates and MDDs (CPI)

Parameter	Models			
	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
Variances of shocks to SV processes				
$\sigma_\eta^2$ (Trend SV)	0.028 [0.015, 0.039]	0.010 [0.008, 0.012]	0.054 [0.042, 0.070]	0.023 [0.018, 0.030]
$\sigma_v^2$ (Gap SV)	0.052 [0.037, 0.089]	0.024 [0.020, 0.031]	0.095 [0.047, 0.129]	0.080 [0.064, 0.102]
Persistence of inflation gap				
$\theta$	0.315 [0.165, 0.463]	0.330 [0.207, 0.444]	-	-
$\sigma_\phi^2$	-	-	0.007 [0.005, 0.008]	0.012 [0.010, 0.015]
Forecast stickiness				
$\lambda$	-	0.163 [0.116, 0.218]	-	0.392 [0.327, 0.460]
$\sigma_\kappa^2$	0.003 [0.001, 0.003]	-	0.003 [0.002, 0.003]	-
Measurement error variances				
$\sigma_{\xi, \pi}^2$	0.168 [0.131, 0.239]	4.710 [3.882, 5.717]	1.032 [0.839, 1.349]	0.130 [0.107, 0.162]
$\sigma_{\xi, 1}^2$	0.531 [0.437, 0.638]	0.011 [0.009, 0.013]	0.403 [0.335, 0.488]	0.577 [0.490, 0.710]
$\sigma_{\xi, 2}^2$	0.071 [0.059, 0.085]	0.011 [0.009, 0.013]	0.054 [0.044, 0.067]	0.071 [0.059, 0.085]
$\sigma_{\xi, 3}^2$	0.027 [0.022, 0.034]	0.011 [0.009, 0.013]	0.029 [0.023, 0.036]	0.026 [0.022, 0.031]
$\sigma_{\xi, 4}^2$	0.024 [0.020, 0.030]	0.011 [0.009, 0.013]	0.025 [0.021, 0.031]	0.024 [0.020, 0.029]
$\sigma_{\xi, 5}^2$	0.030 [0.025, 0.037]	0.011 [0.009, 0.013]	0.027 [0.023, 0.033]	0.031 [0.025, 0.037]
$\ln \text{MDD}(\mathcal{M}_i   y^T)$	-447.019 (0.144)	-447.775 (0.139)	-447.582 (0.360)	-452.472 (0.249)


Note: The table contains posterior moments and log MDDs for the state space models  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$  based on  $M = 100,000$  particles and the full data sample. The main entry for every static parameter reports its posterior median with five and 95 percent quantiles in brackets below. Log MDDs for model  $i$  are denoted  $\ln \text{MDD}(\mathcal{M}_i | y^T)$  and computed using equation (17) of the paper. The reported values are the average estimates obtained from 250 repetitions of the particle learning filter, and the associated numerical standard errors appear in parentheses below each estimate.

Figure R.52:  $\mathcal{M}_0$ : Trend and Gap Inflation



Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_0$ . The vertical dotted bands denote NBER dated recessions in the four charts.

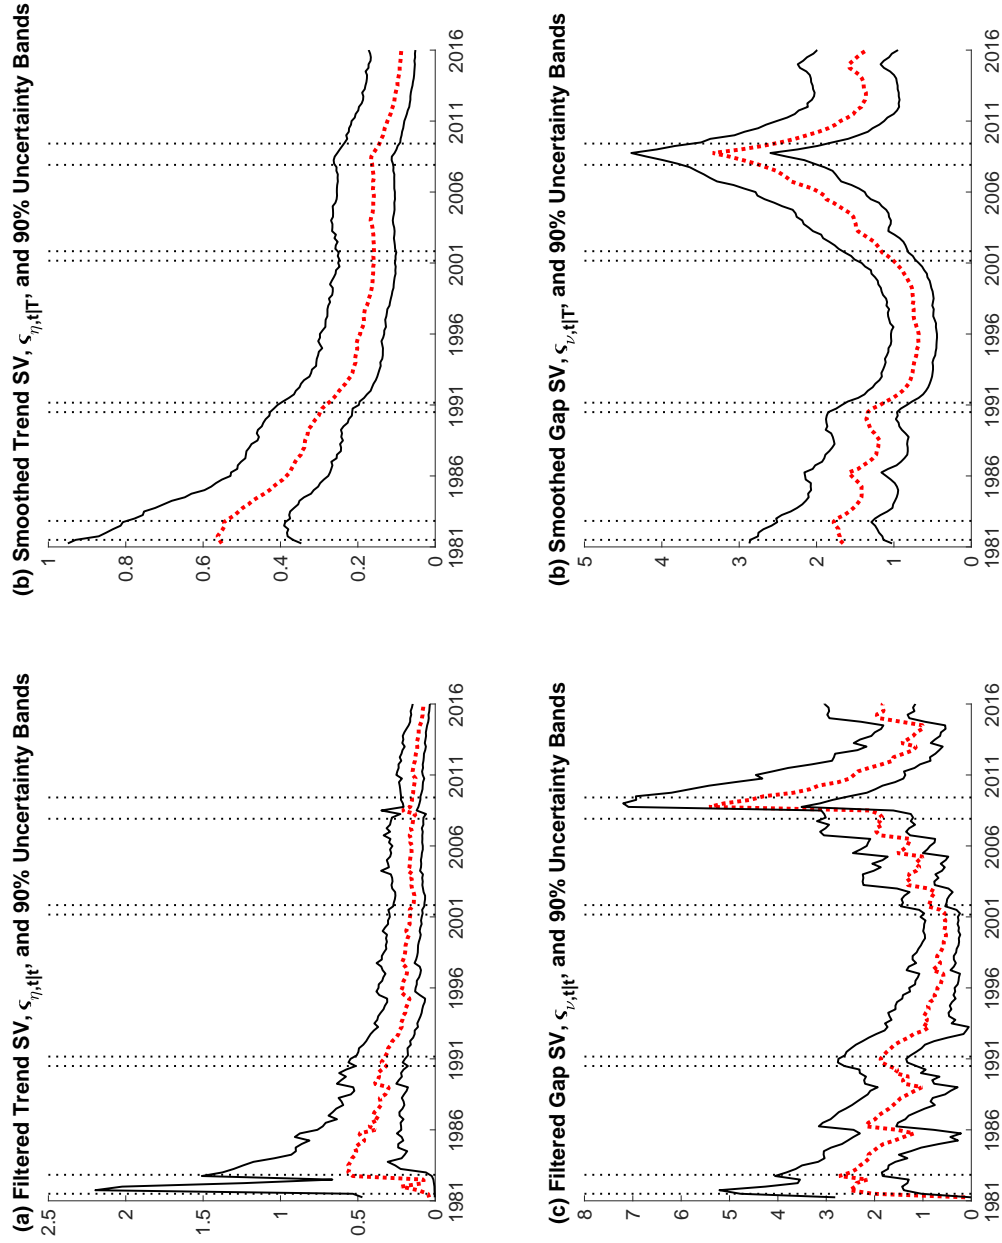
Figure R.53:  $\mathcal{M}_0$ : Static Parameters



Note: Posterior quantiles of particle-learning estimates. Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_0$ . Dotted vertical lines denote NBER recession peaks and troughs.




Figure R.54:  $\mathcal{M}_0$ : Stochastic Volatility in Trend and Gap Inflation



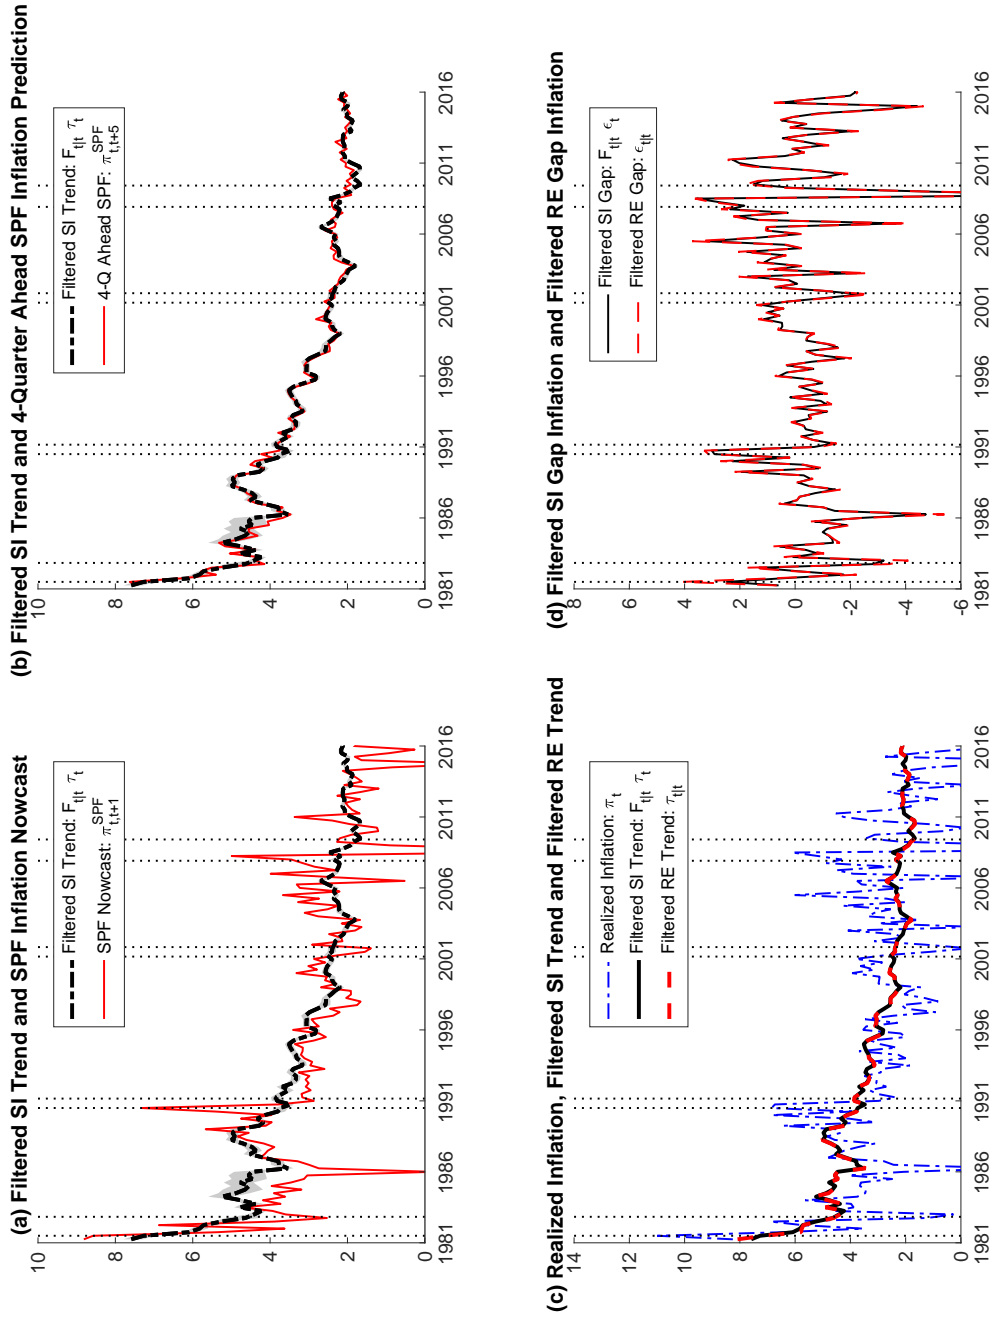
Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_0$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.55:  $\mathcal{M}_0$ : Time-Varying SI Parameter




Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{fit}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands delineated by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{f|T} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_0$ .

Figure R.56:  $\mathcal{M}_1$ : Trend and Gap Inflation



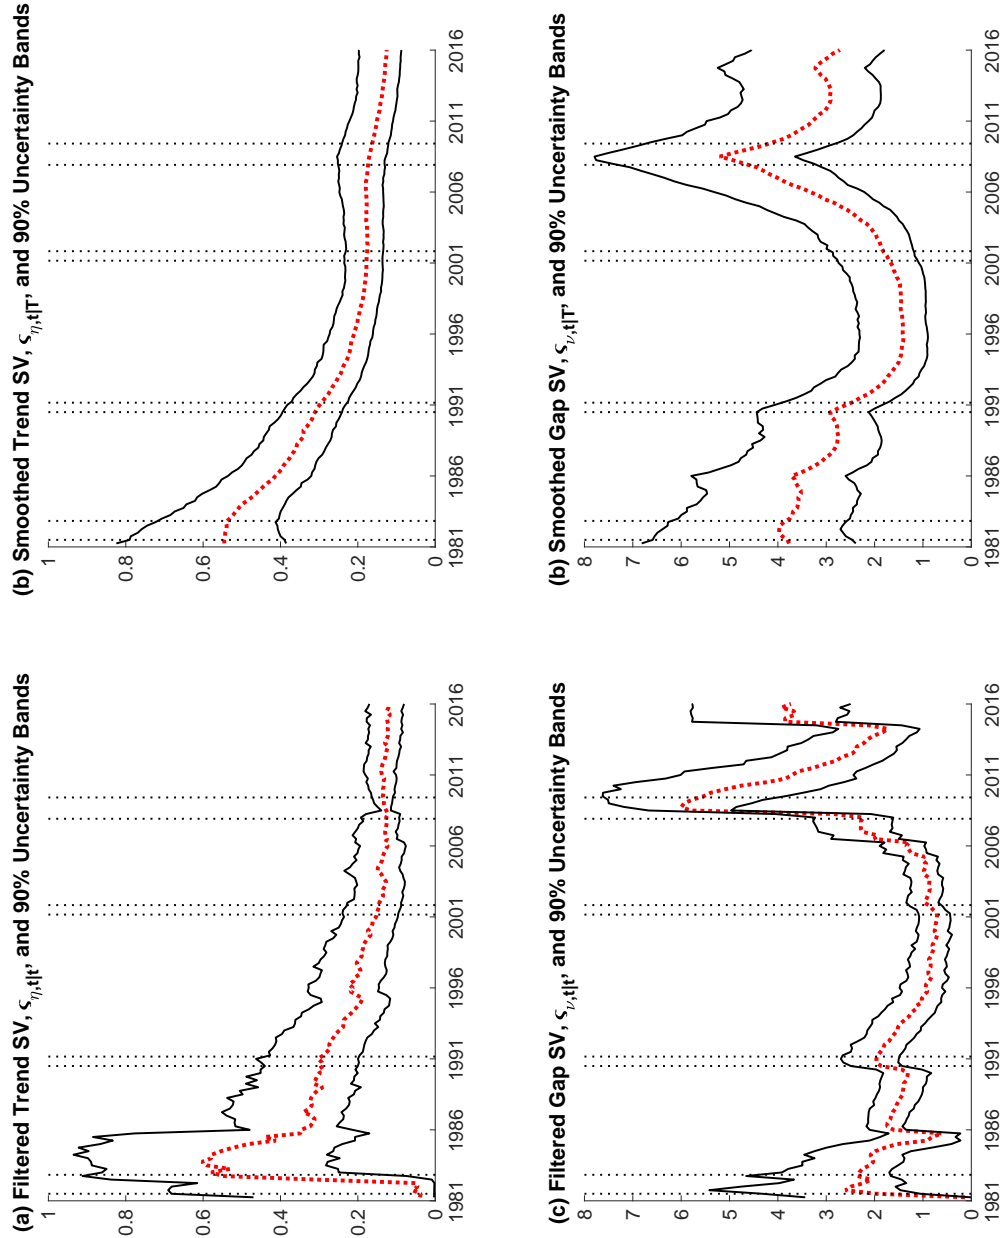
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_1$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.57:  $\mathcal{M}_1$ : Static Parameters



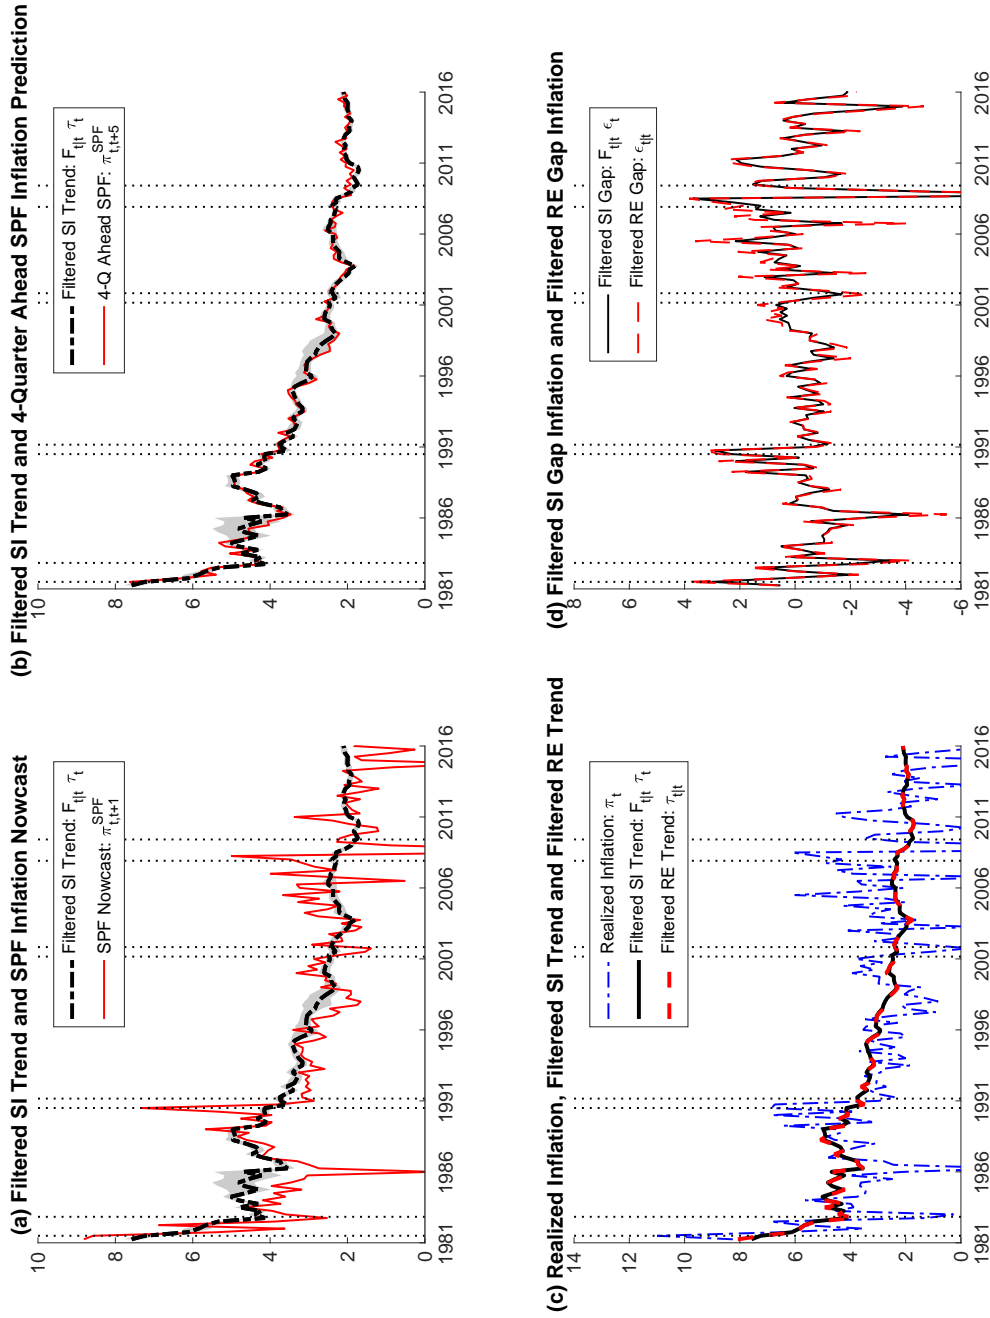
Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_1$ . Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.58:  $\mathcal{M}_1$ : Stochastic Volatility in Trend and Gap Inflation



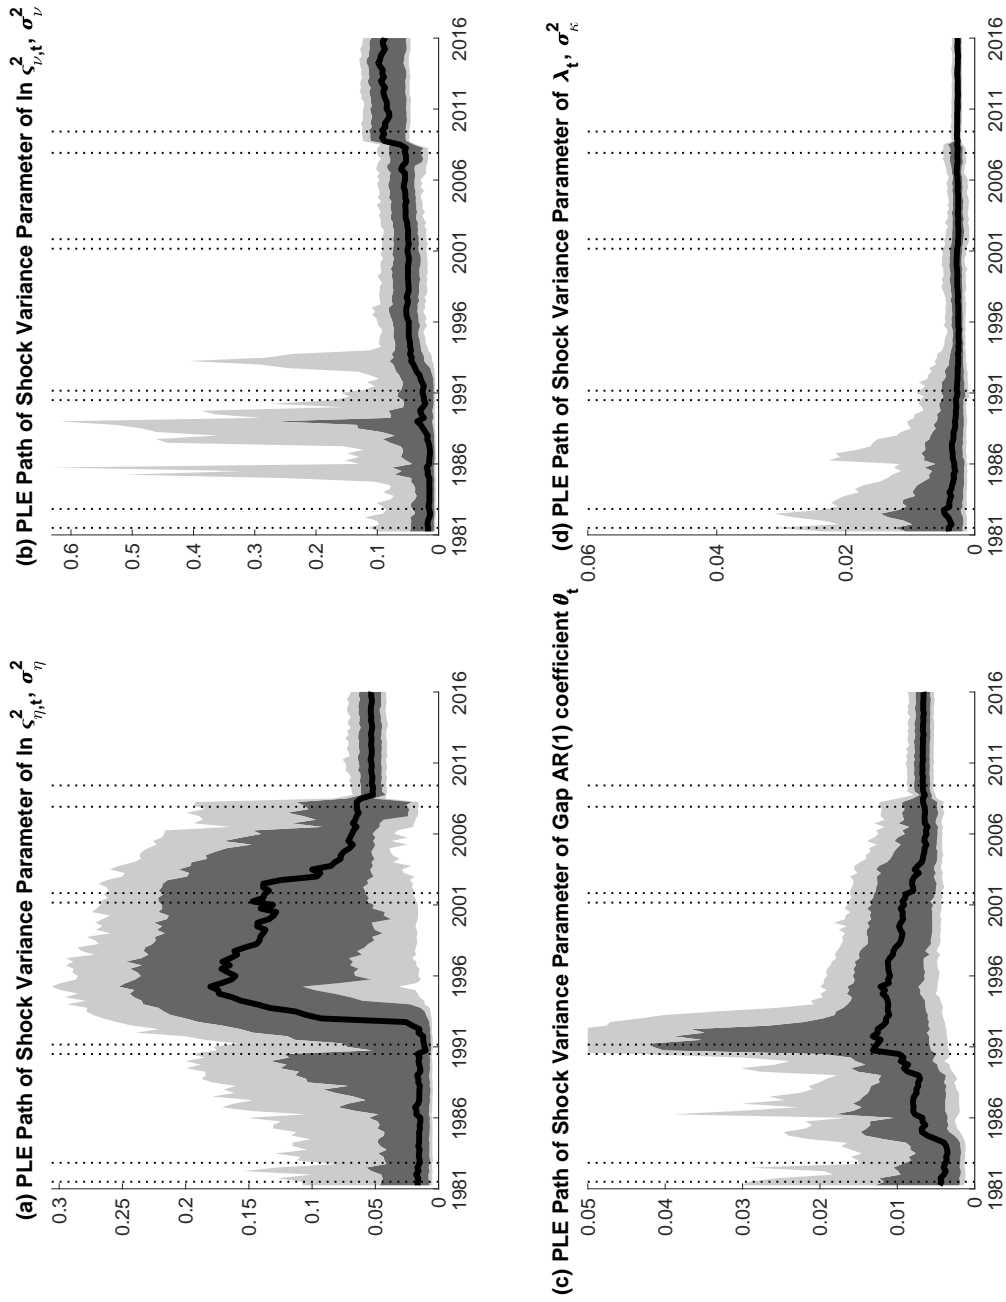
Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_1$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.59:  $\mathcal{M}_2$ : Trend and Gap Inflation



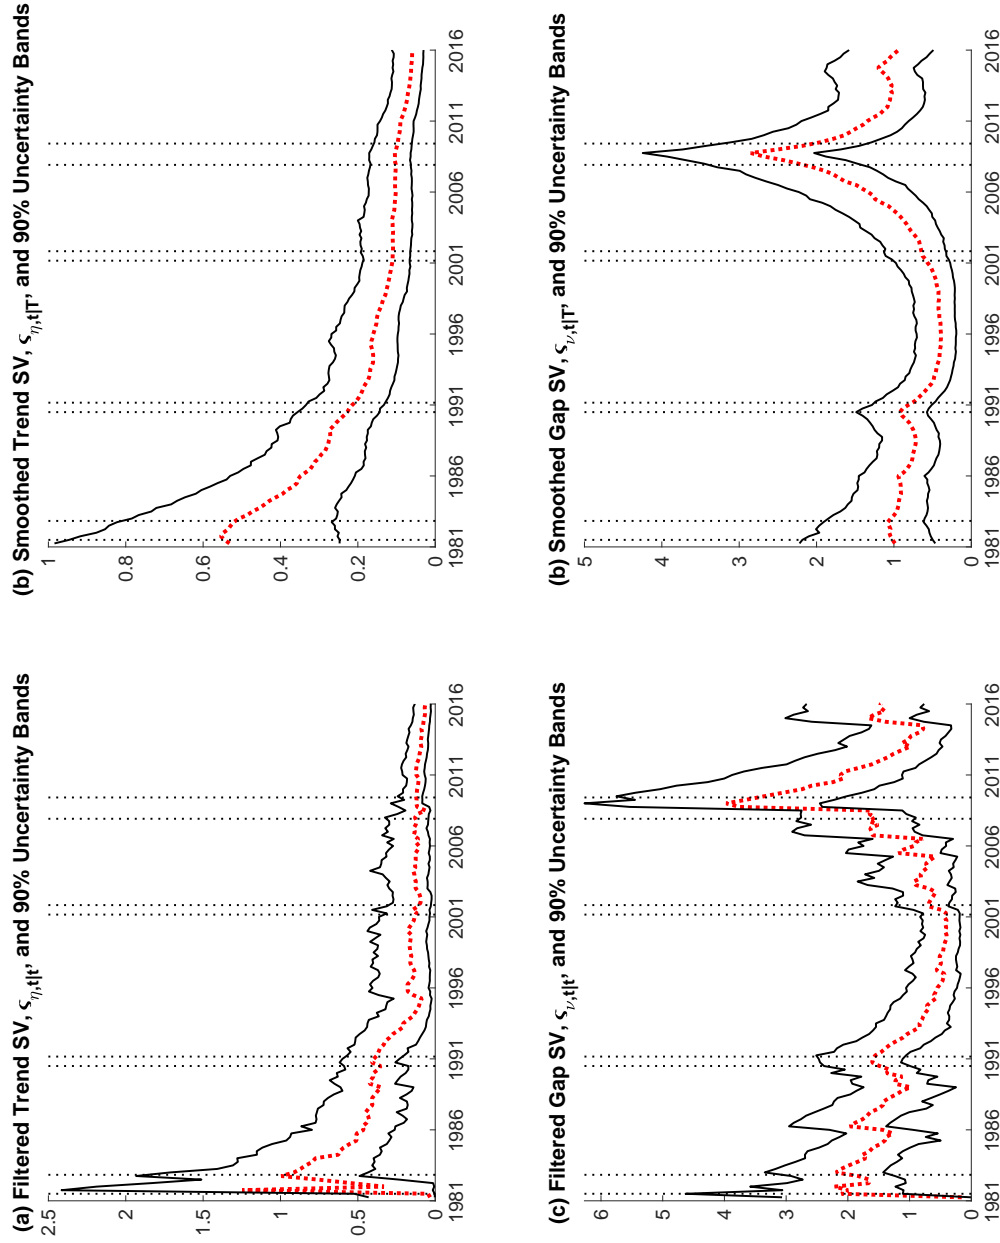
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_2$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.60:  $\mathcal{M}_2$ : Static Parameters



Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_2$ . Dotted vertical lines denote NBER recession peaks and troughs.


Figure R.61:  $\mathcal{M}_2$ : Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_2$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.




Figure R.62:  $\mathcal{M}_2$ : Time-Varying SI Parameter



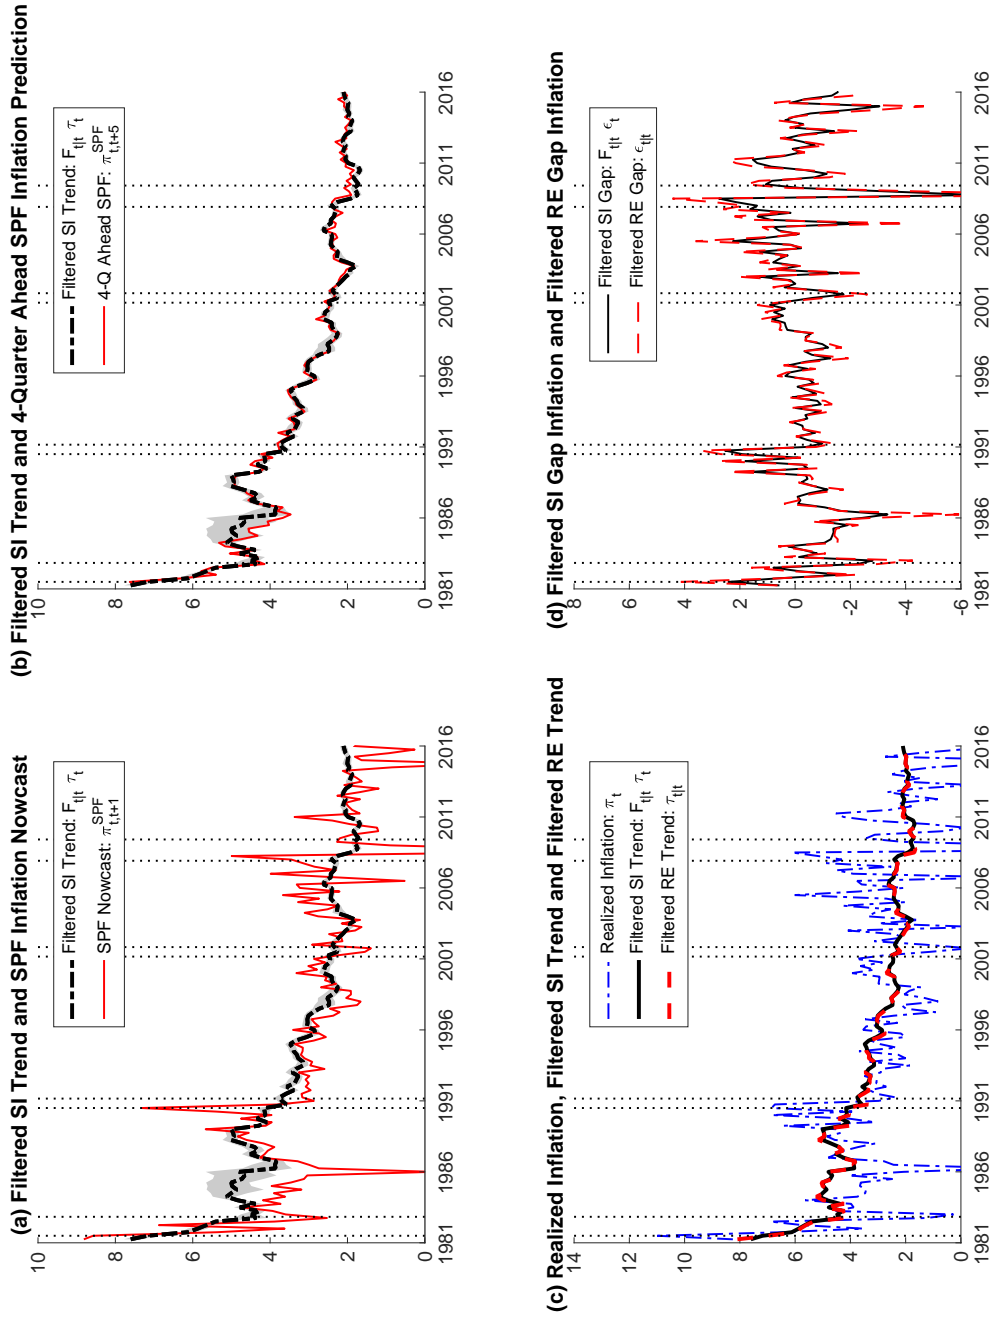
Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{fit}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{f|T} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_2$ .

Figure R.63:  $\mathcal{M}_2$ : Time-Varying AR Coefficient in the Inflation Gap Process



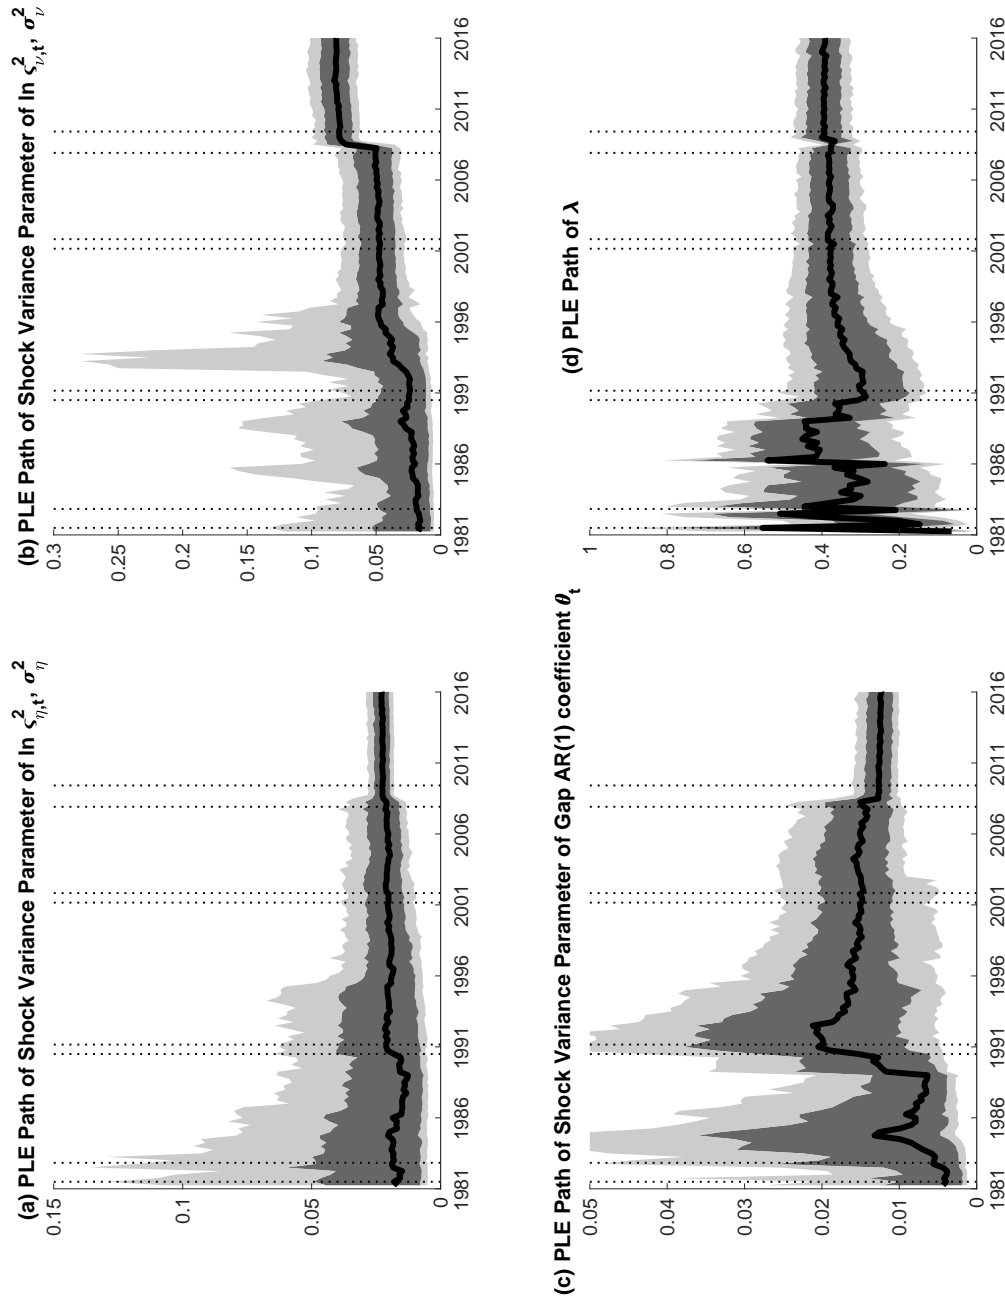
Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{ft}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{ftT}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_2$ .

Figure R.64:  $\mathcal{M}_3$ : Trend and Gap Inflation



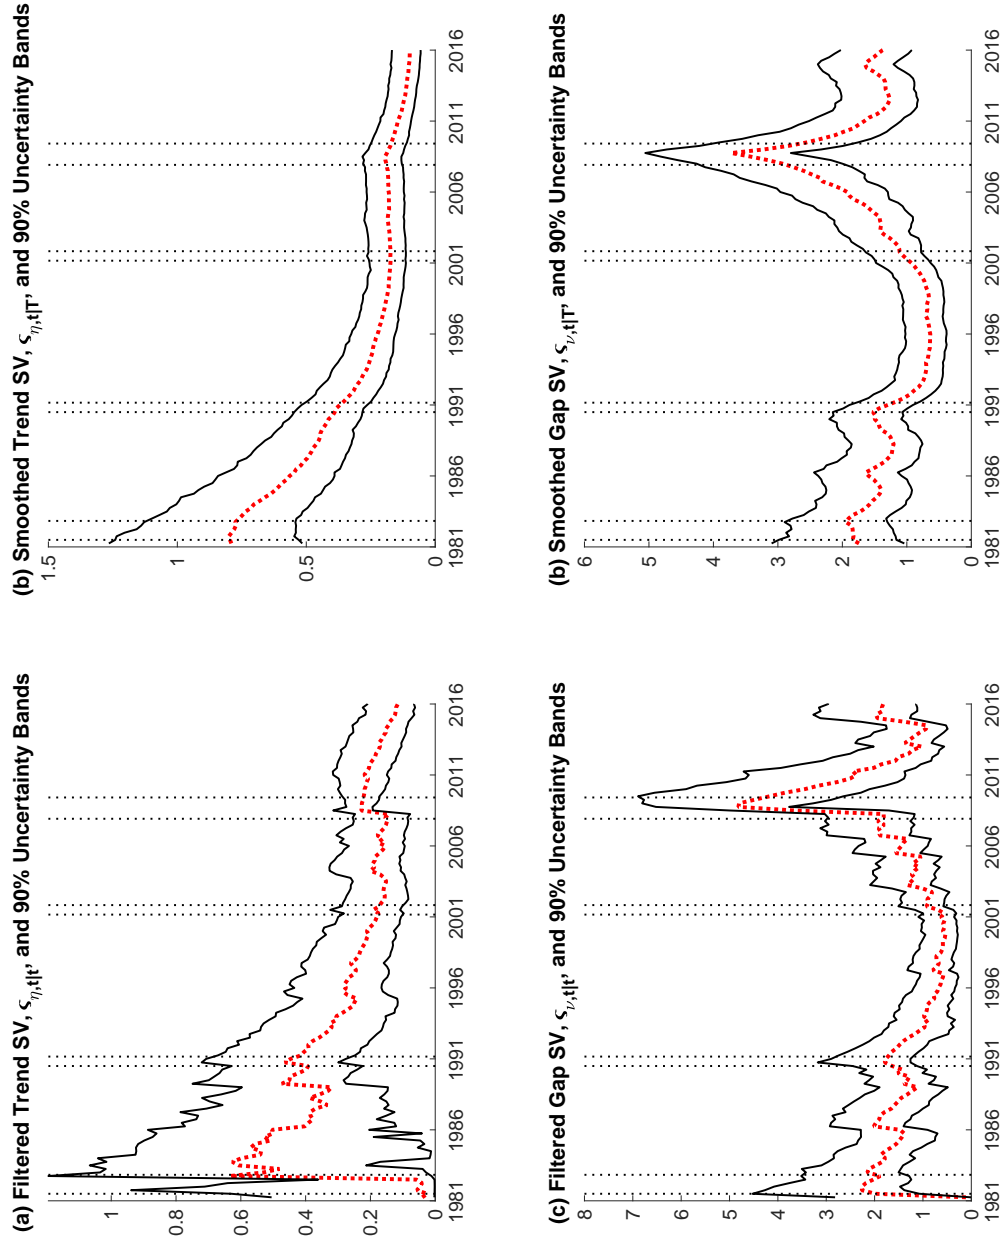
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t} \tau_t$  generated by Model  $\mathcal{M}_3$ . The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.65:  $\mathcal{M}_3$ : Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_3$ . Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.66:  $\mathcal{M}_3$ : Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_3$ , are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.67:  $\mathcal{M}_3$ : Time-Varying AR Coefficient in the Inflation Gap Process



Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{t|t}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{t|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_3$ .

## R.7.2 CPI results for case without noise component in inflation

This section reports results for alternative versions of each of our four model variants. These alternative state space models shut off the irregular component (or noise) in the inflation equation. In this case, the inflation process reduces to

$$\pi_t = \tau_t + \varepsilon_t, \tag{R.23}$$

with the remaining elements of each model, including the specification of priors, unchanged. Please note that, as in the baseline specification, measurement error is retained in the equations mapping SI forecasts into SPF forecast data.

Table R.5 reports estimates of static parameters as well as log MDDs for each of the four model variants when noise in the inflation equation has been removed. In comparison to results from the baseline specification (with noise in inflation), as reported in Table 5 of the paper, the logMDDs are generally lower (and significantly so), providing strong evidence against removing the noise component from the inflation equation. However, when comparing these alternative model variants against each other, the  $\mathcal{M}_2$  model with time-varying  $\theta_t$  and  $\lambda_t$  continues to attract the highest log MDD (as reported in the paper for the case with noise in inflation).

Detailed results are shown separately for each model variant in the following figures: Figures R.19-R.22 report results for  $\mathcal{M}_0$ , with time-varying  $\lambda_t$  but constant  $\theta$ . Figures R.23-R.25 show results for  $\mathcal{M}_1$ , where both  $\lambda$  and  $\theta$  are assumed to be constant. Figures R.26-R.30 provide results for  $\mathcal{M}_2$ , where both  $\lambda_t$  and  $\theta_t$  are time-varying. Finally, Figures R.31-R.34 show results for  $\mathcal{M}_3$ , where  $\theta_t$  is time-varying but  $\lambda$  is assumed constant.

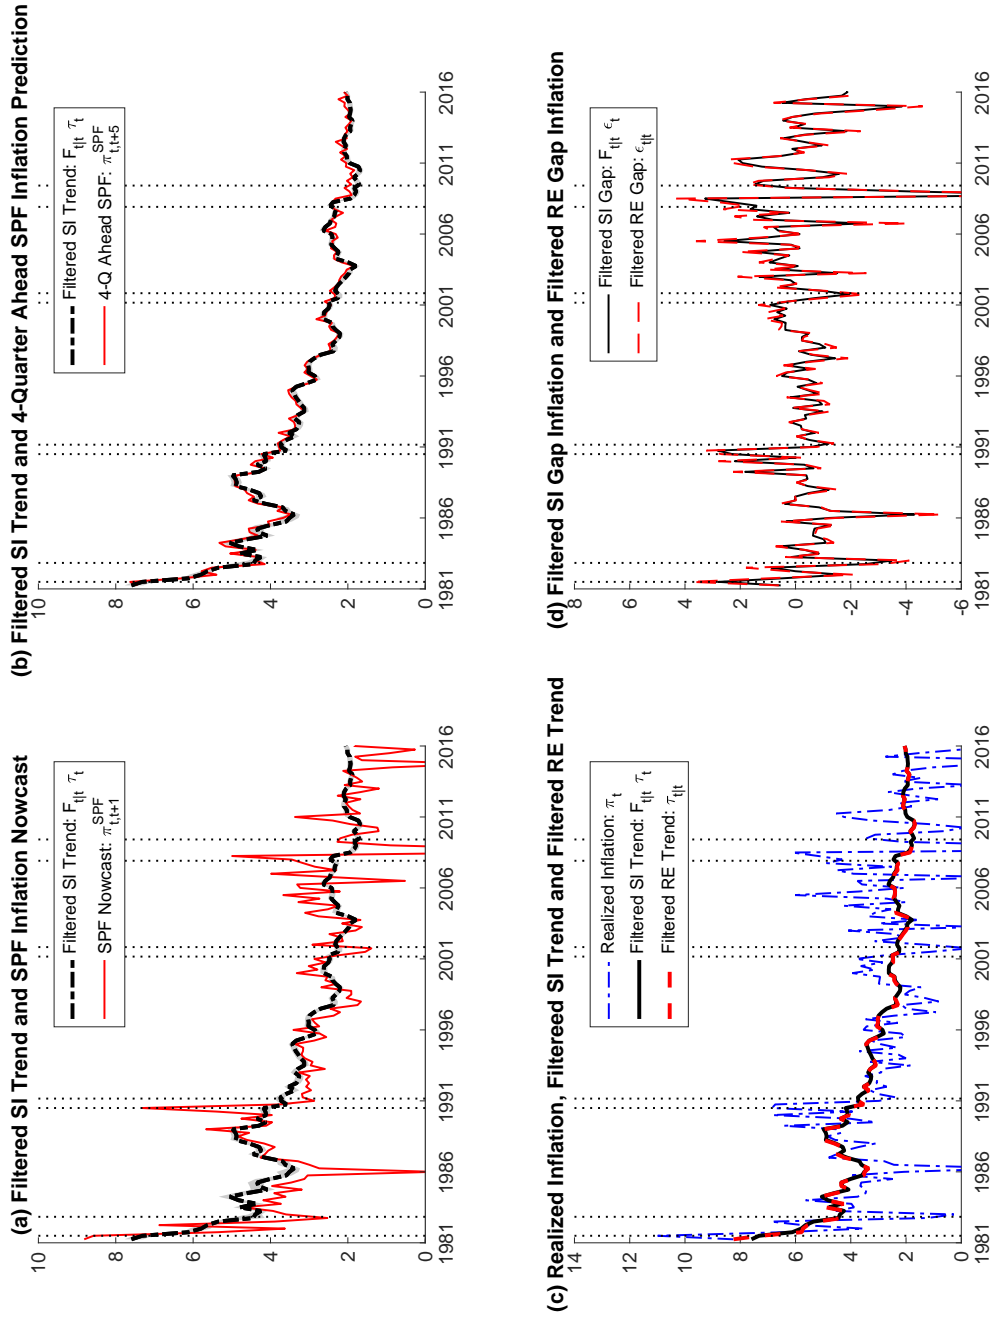
Table R.5: Parameter Estimates and MDDs w/o noise in inflation (CPI)

Parameter	Models			
	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
Variances of shocks to SV processes				
$\sigma_\eta^2$ (Trend SV)	0.054 [0.023, 0.082]	0.009 [0.007, 0.012]	0.011 [0.008, 0.016]	0.047 [0.028, 0.068]
$\sigma_v^2$ (Gap SV)	0.036 [0.025, 0.091]	0.014 [0.011, 0.017]	0.228 [0.067, 0.302]	0.036 [0.026, 0.055]
Persistence of inflation gap				
$\theta$	0.226 [0.088, 0.373]	0.176 [0.043, 0.308]	-	-
$\sigma_\phi^2$	-	-	0.007 [0.004, 0.009]	0.008 [0.004, 0.011]
Forecast stickiness				
$\lambda$	-	0.129 [0.087, 0.176]	-	0.333 [0.265, 0.406]
$\sigma_\kappa^2$	0.001 [0.001, 0.003]	-	0.003 [0.001, 0.004]	-
Measurement error variances				
$\sigma_{\xi, \pi}^2$	-	-	-	-
$\sigma_{\xi, 1}^2$	0.508 [0.421, 0.626]	0.011 [0.009, 0.013]	0.592 [0.495, 0.722]	0.615 [0.517, 0.746]
$\sigma_{\xi, 2}^2$	0.071 [0.058, 0.086]	0.011 [0.009, 0.013]	0.089 [0.074, 0.106]	0.077 [0.063, 0.094]
$\sigma_{\xi, 3}^2$	0.029 [0.023, 0.035]	0.011 [0.009, 0.013]	0.033 [0.026, 0.040]	0.028 [0.023, 0.034]
$\sigma_{\xi, 4}^2$	0.027 [0.022, 0.033]	0.011 [0.009, 0.013]	0.025 [0.021, 0.030]	0.027 [0.022, 0.033]
$\sigma_{\xi, 5}^2$	0.035 [0.029, 0.042]	0.011 [0.009, 0.013]	0.028 [0.023, 0.035]	0.033 [0.027, 0.041]
$\ln \text{MDD}(\mathcal{M}_i   y^T)$	-448.830 (0.109)	-450.470 (0.130)	-462.330 (0.137)	-461.666 (0.142)

Note: The table contains posterior moments and log MDDs for the state space models  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$  based on  $M = 100,000$  particles and the full data sample. The main entry for every static parameter reports its posterior median with five and 95 percent quantiles in brackets below. Log MDDs for model  $i$  are denoted  $\ln \text{MDD}(\mathcal{M}_i | y^T)$  and computed using equation (17) of the paper. The reported values are the average estimates obtained from 250 repetitions of the particle learning filter, and the associated numerical standard errors appear in parentheses below each estimate.

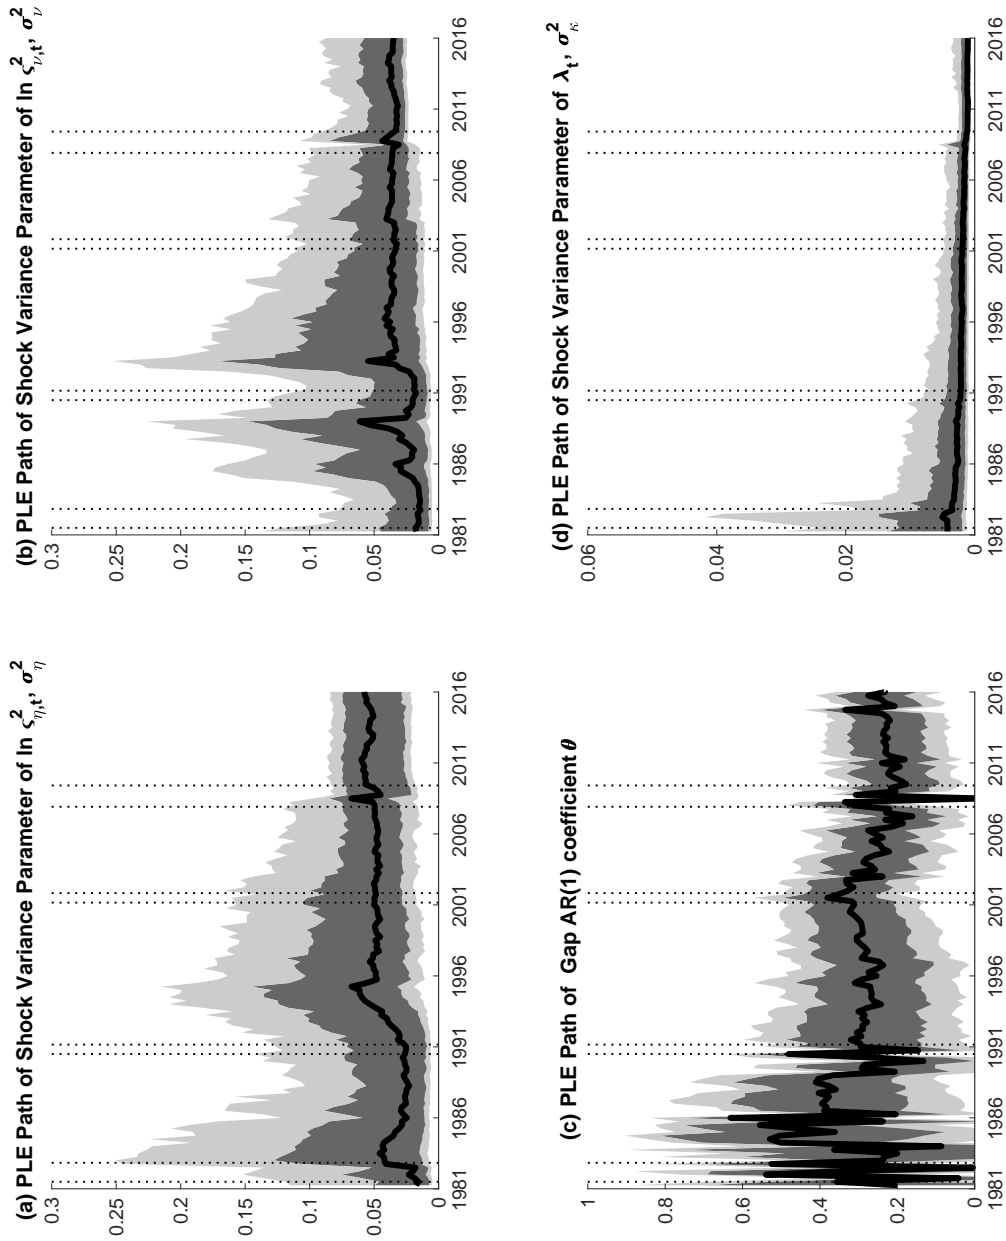


Figure R.68:  $\mathcal{M}_0$  (w/o noise): Trend and Gap Inflation



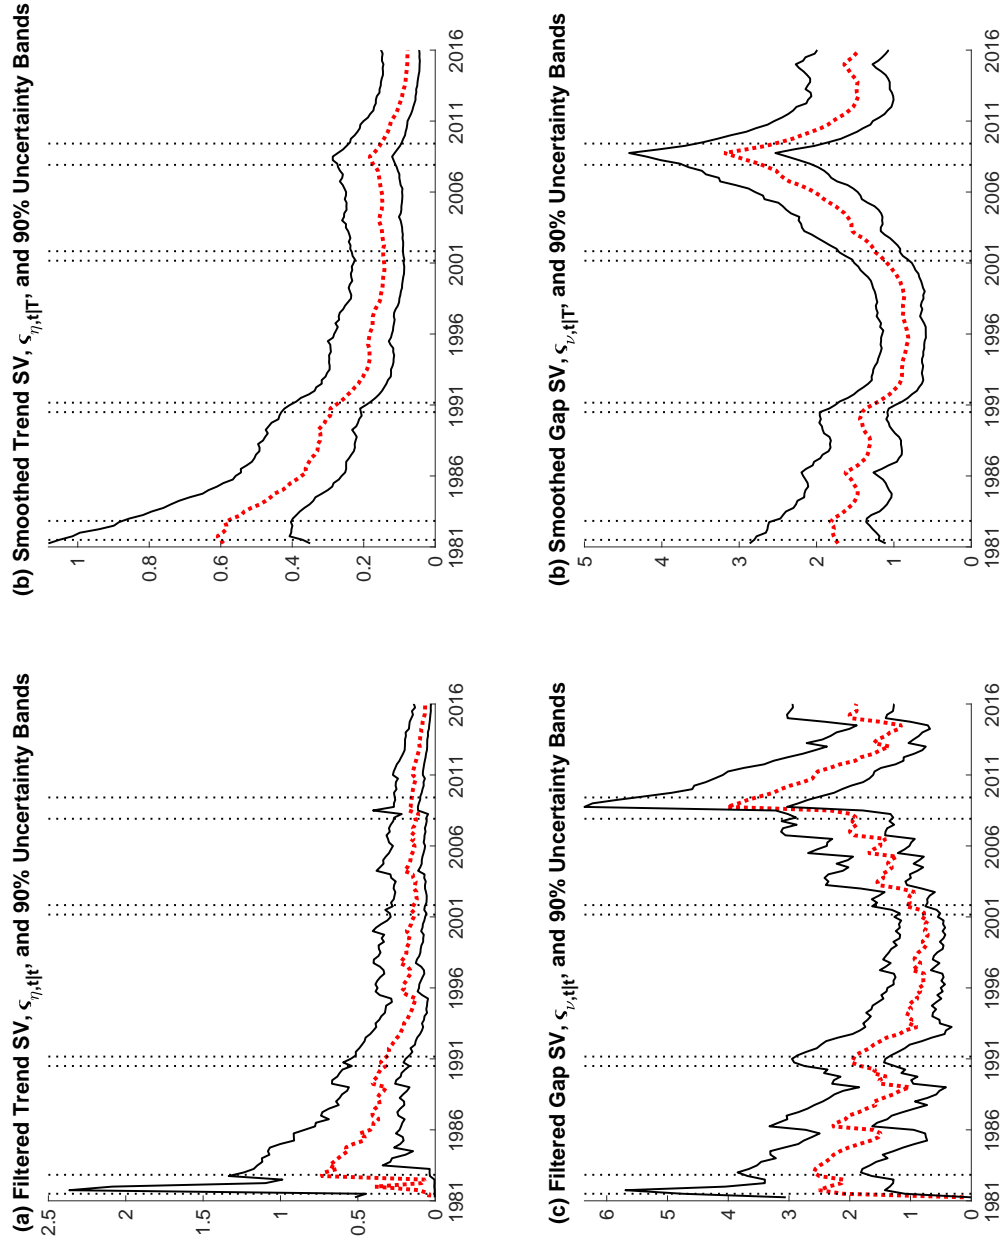
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{t|t}\tau_t$  generated by Model  $\mathcal{M}_0$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.69:  $\mathcal{M}_0$  (w/o noise): Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_0$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.70:  $\mathcal{M}_0$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



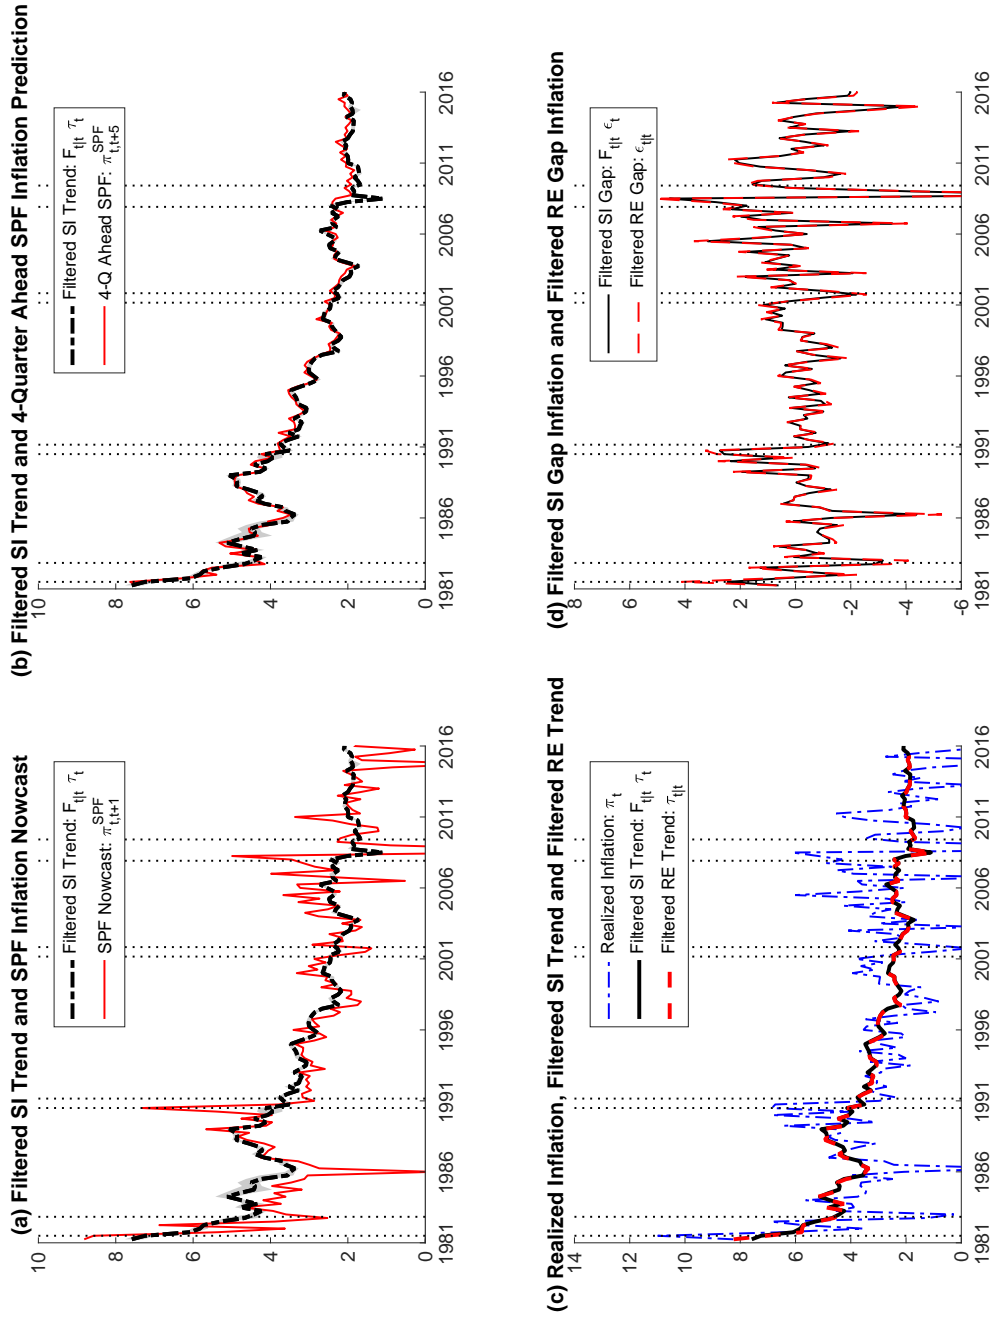
Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_0$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.71:  $\mathcal{M}_0$  (w/o noise): Time-Varying SI Parameter




Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{fit}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands delineated by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{f|T} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_0$  (w/o noise).

Figure R.72:  $\mathcal{M}_1$  (w/o noise): Trend and Gap Inflation



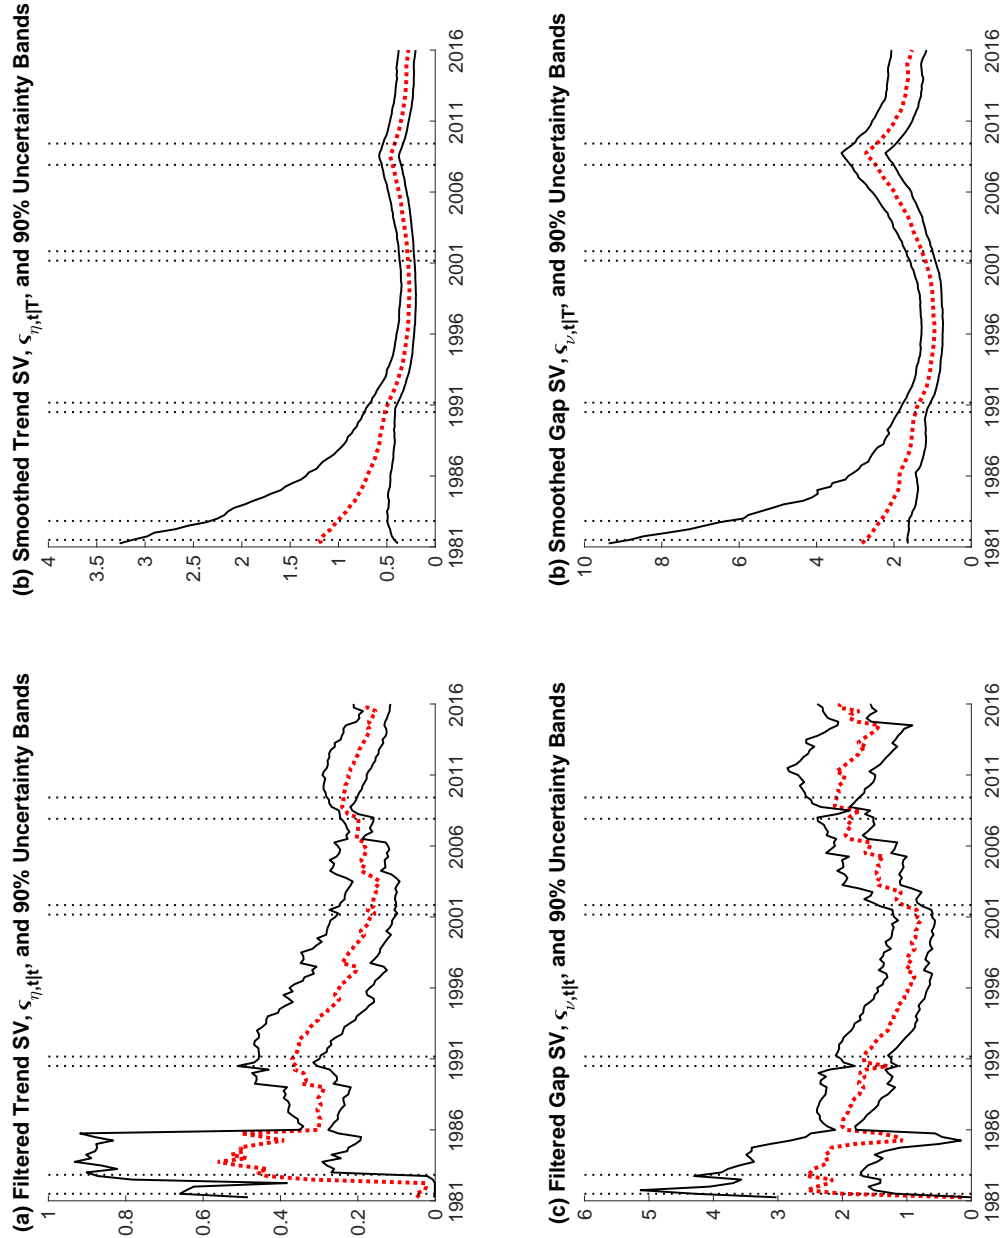
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{\text{fit}}\tau_t$  generated by Model  $\mathcal{M}_1$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.73:  $\mathcal{M}_1$  (w/o noise): Static Parameters



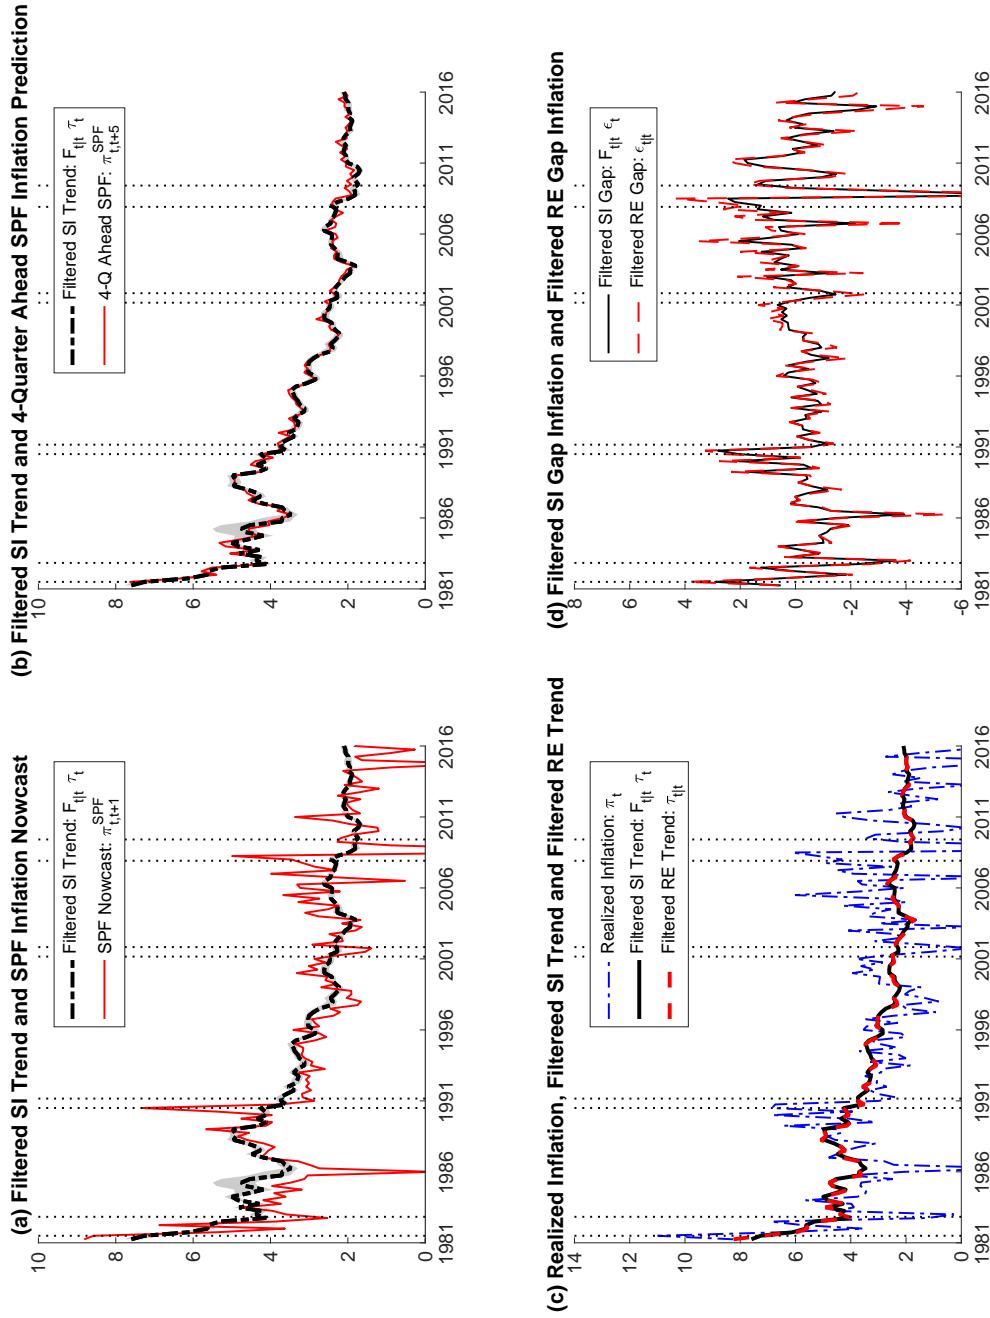
Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_1$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.74:  $\mathcal{M}_1$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_1$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

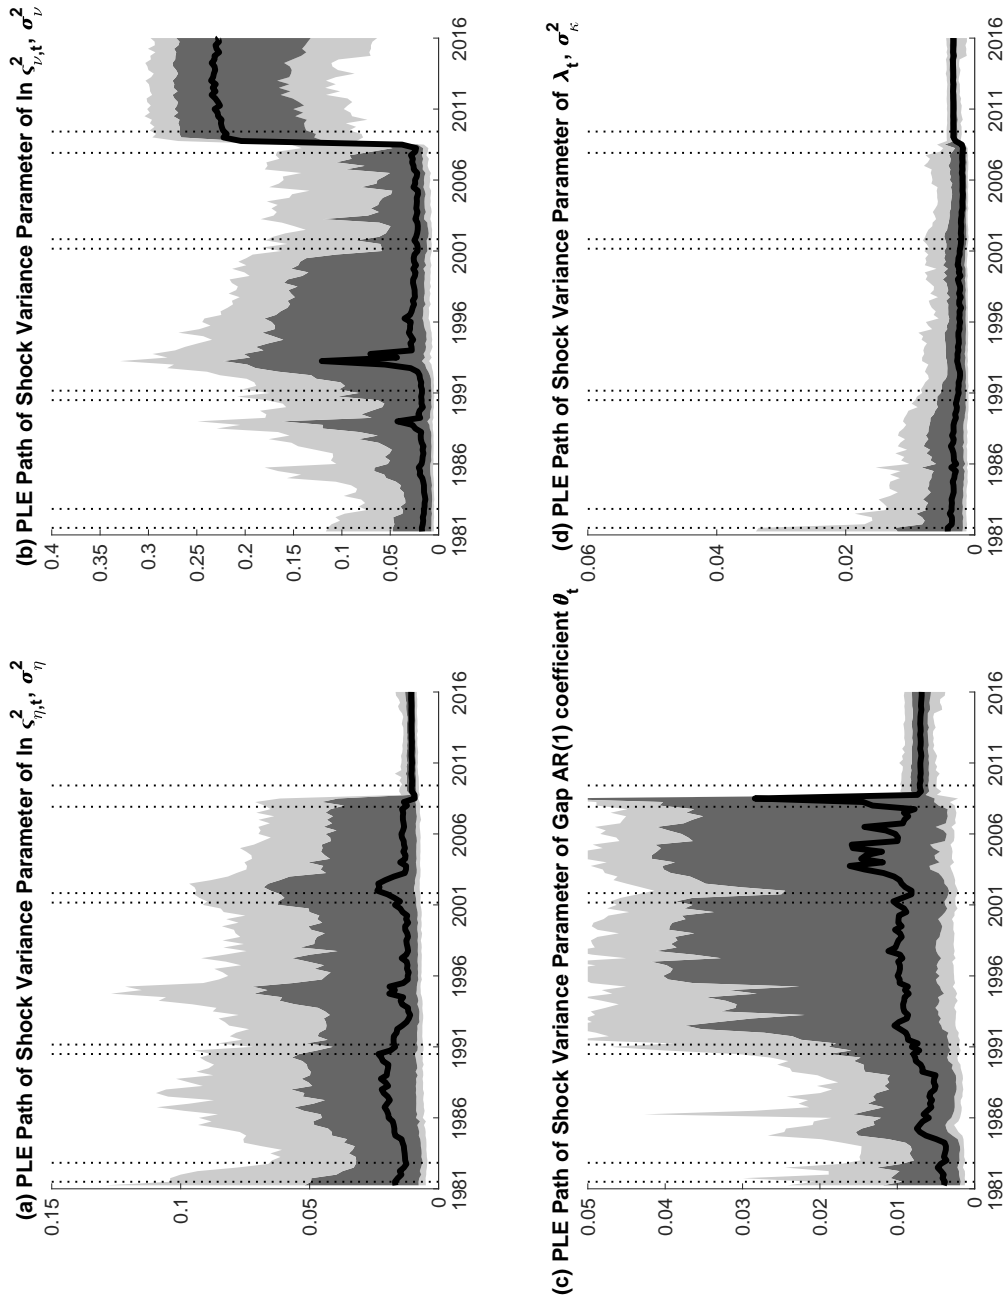
Figure R.75:  $\mathcal{M}_2$  (w/o noise): Trend and Gap Inflation



Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{\text{fit}}\tau_t$  generated by Model  $\mathcal{M}_2$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

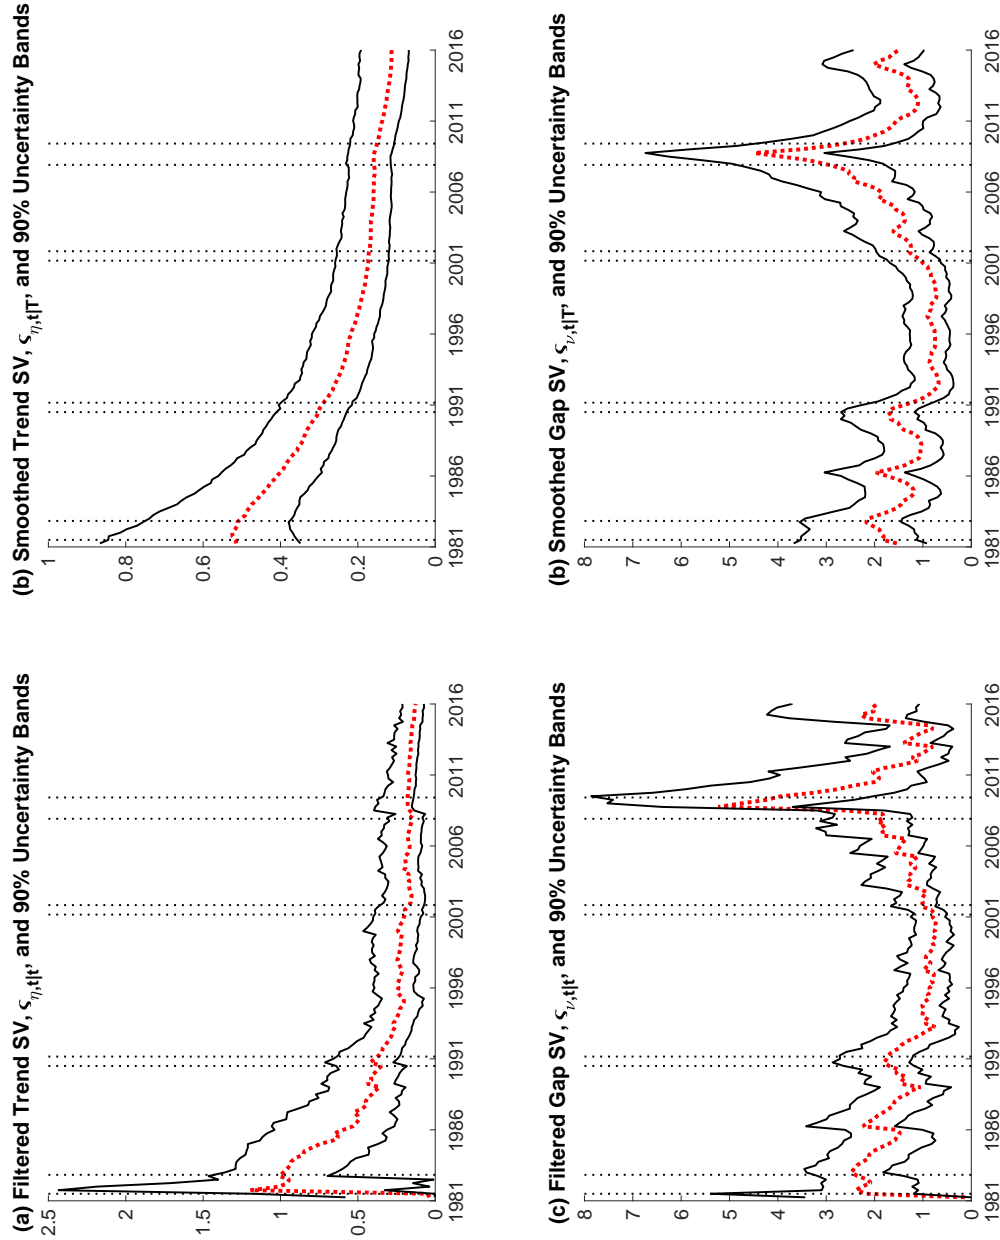


Figure R.76:  $\mathcal{M}_2$  (w/o noise): Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_2$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.77:  $\mathcal{M}_2$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation




Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_2$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.78:  $\mathcal{M}_2$  (w/o noise): Time-Varying SI Parameter



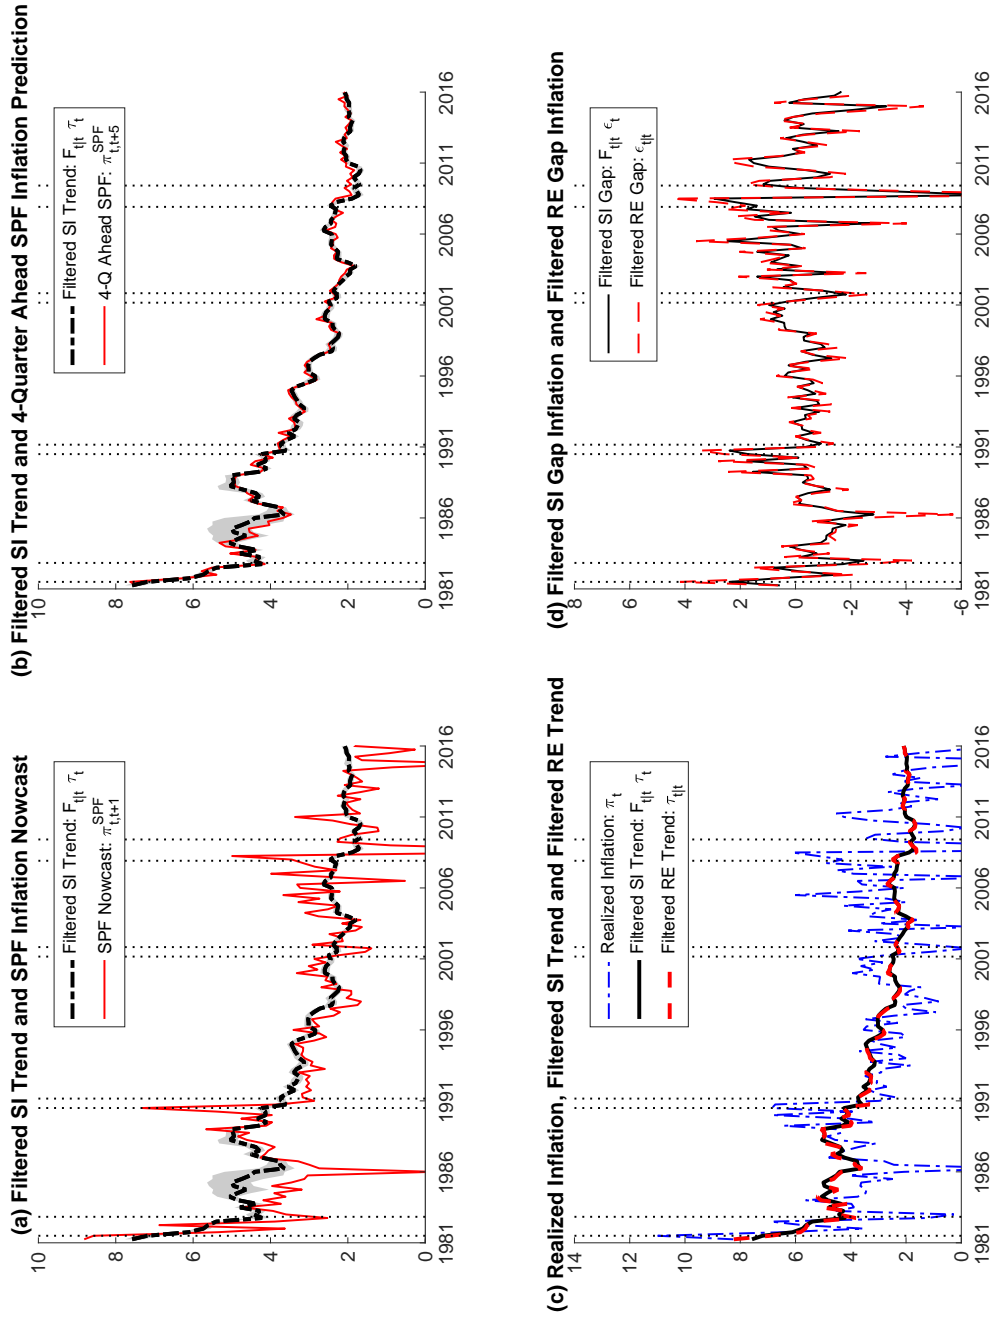
Note: In the top panel, dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{fit}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. The bottom panel displays estimated differences  $\lambda_{f|T} - \lambda_{0|T}$  of the smoothed estimates with corresponding 68% (90%) uncertainty bands shown as dark (light) gray areas. All estimates generated from model  $\mathcal{M}_2$  (w/o noise).

Figure R.79:  $\mathcal{M}_2$  (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process



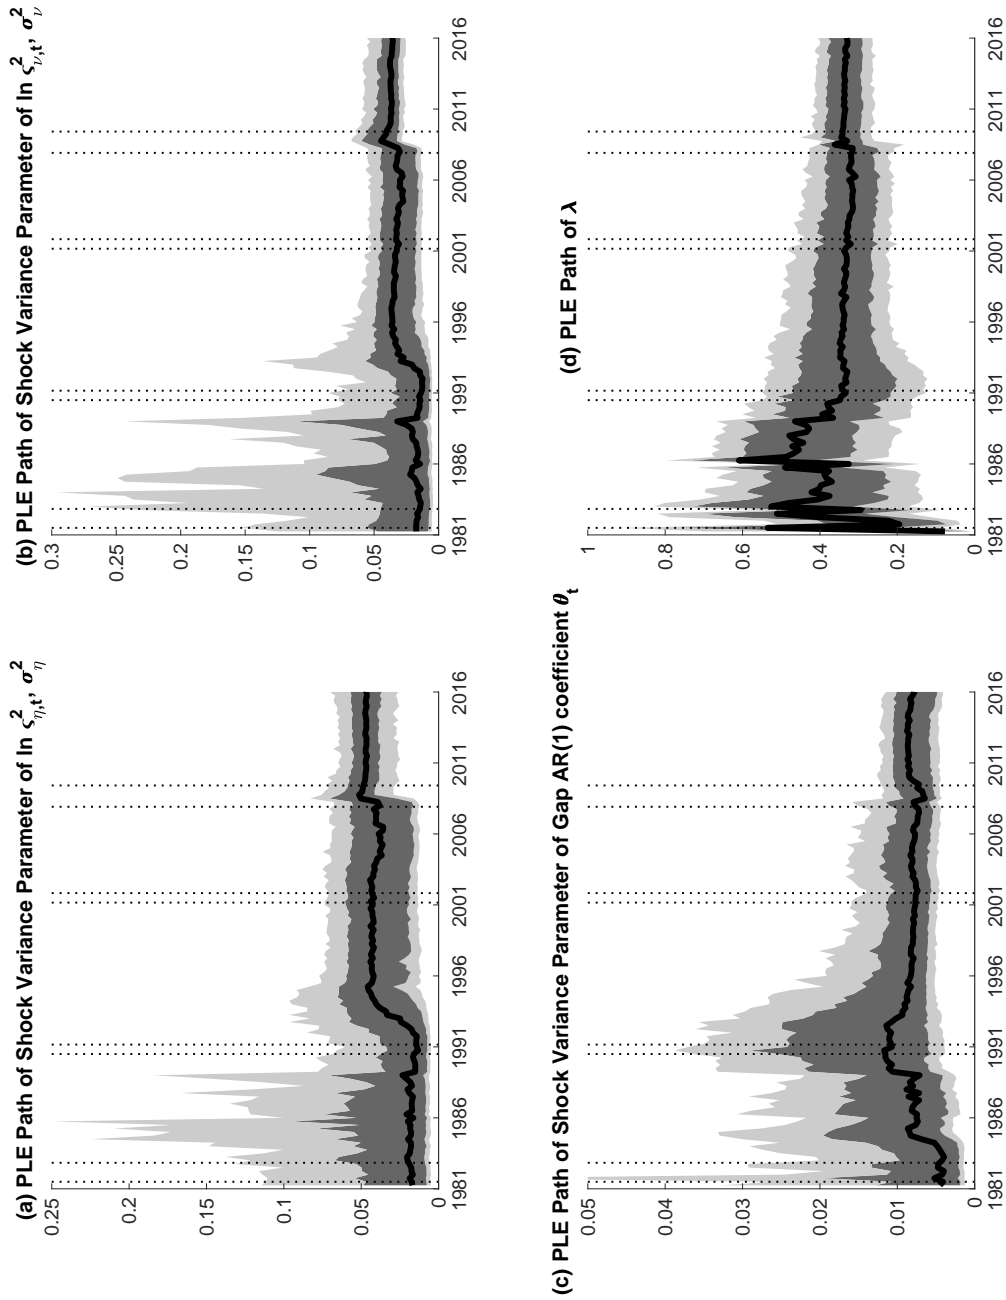
Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{ft}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{ftT}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_2$  (w/o noise).

Figure R.80:  $\mathcal{M}_3$  (w/o noise): Trend and Gap Inflation



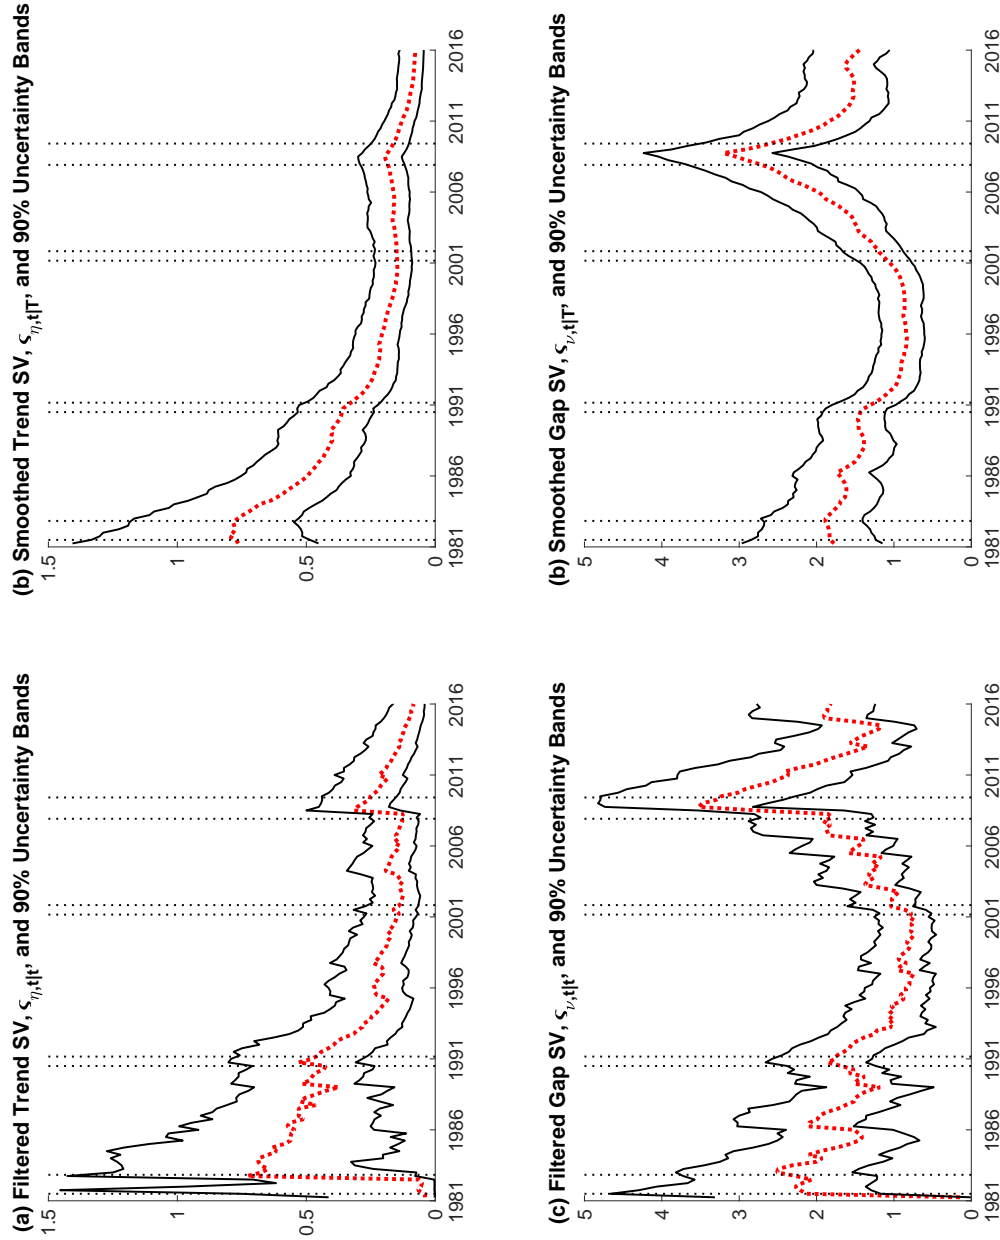
Note: The top row of charts contains light gray shaded areas that represent 68 percent uncertain bands around estimates of filtered SI trend inflation,  $F_{\text{fit}}\tau_t$  generated by Model  $\mathcal{M}_3$  (w/o noise). The vertical dotted bands denote NBER dated recessions in the four charts.

Figure R.81:  $\mathcal{M}_3$  (w/o noise): Static Parameters




Note: Posterior quantiles of particle-learning estimates (PLE). Solid line depicts median, dark and light shaded areas correspond to 68% and 90% uncertainty bands, respectively, as estimated from model  $\mathcal{M}_3$  (w/o noise). Dotted vertical lines denote NBER recession peaks and troughs.

Figure R.82:  $\mathcal{M}_3$  (w/o noise): Stochastic Volatility in Trend and Gap Inflation



Note: The solid thin (black) lines around estimates of filtered and smoothed SV in shocks to trend and gap inflation, estimated from model  $\mathcal{M}_3$  (w/o noise), are lower and upper bounds on 90% uncertainty bands. The four plots contain vertical dotted bands that denote NBER dated recessions.

Figure R.83:  $\mathcal{M}_3$  (w/o noise): Time-Varying AR Coefficient in the Inflation Gap Process



Note: Dark (light) gray areas are 68% (90%) uncertainty bands around filtered estimates  $\lambda_{f|t}$  depicted by the dashed (black) line. Solid thin (red) lines show smoothed estimates  $\lambda_{f|T}$  surrounded by 90% uncertainty bands that are depicted by the dot-dashed (red) lines. All estimates generated from model  $\mathcal{M}_3$  (w/o noise).



### R.7.3 Log MDD for different numbers of particles (CPI)

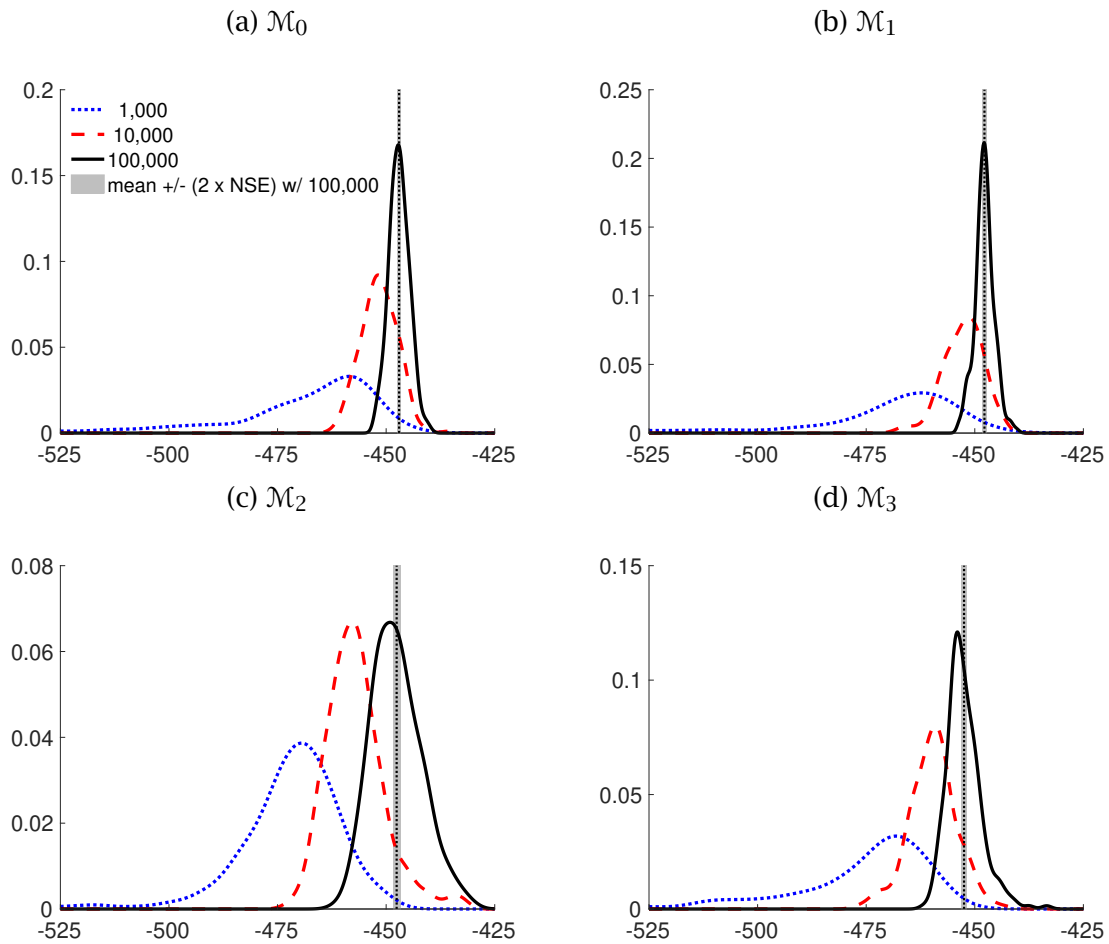
Figure R.84 displays the distribution of log MDD estimates generated by  $N = 250$  repetitions of the particle learning filter applied to each of our model variants. For each model, distributions of log MDD estimates are generated for  $M = 100,000$ ,  $M = 10,000$ , and  $M = 1,000$  particles. These distributions represent uncertainty from the Monte Carlo approximation of the true log MDD associated with each model. We gauge the uncertainty of the simulation-based estimates of the log MDD with numerical standard errors. Denoting the log MDD estimate generated by the  $n$ -th simulation  $\mu_n$ , the numerical standard errors around the average estimate,  $\bar{\mu} = \sum_{n=1}^N \mu_n / N$ , are given by

$$\sigma_{\bar{\mu}} = \sqrt{\frac{\sigma_{\mu}^2}{N}} \quad \text{with} \quad \sigma_{\mu}^2 = \frac{1}{N} \sum_{n=1}^N (\mu_n - \bar{\mu})^2 . \quad (\text{R.24})$$

The use of numerical standard errors for gauging the uncertainty of simulation-based estimates is grounded in the work of Geweke (1989), see also Fuentes-Albero and Melosi (2013), and Herbst and Schorfheide (2014) for applications in the context of log MDD estimates.


As reported in Table 5 of the paper, with  $M = 100,000$ , the numerical standard errors are fairly tight, and the mean log MDD estimates are clearly distinguished across models. As shown in Figure R.84, for choices of the number of particles lower than  $M = 100,000$ , the simulated log MDD distributions display more considerable dispersion, fatter tails and even some skew. Figure R.85 presents similar results for the four model variants when the noise component in the inflation equation is set to zero.

Figure R.84: log MDD for different numbers of particles (CPI)



Note: Distribution of log MDD estimates generated by 250 repetitions of the particle learning filter applied to each model variant. For each model, distributions of log MDD estimates are generated for  $M = 100,000$  (solid, black),  $M = 10,000$  (dashed, red) and  $M = 1,000$  (dotted, blue) particles. In each case, the distributions shown are kernel density estimates obtained from the 250 simulated log MDD values. The horizontal dashed (black) line, displays the average estimate obtained with  $M = 100,000$  particles, and the surrounding shaded area demarcates bands of plus/minus twice the estimated numerical standard error.

Figure R.85: log MDD for different numbers of particles (w/o noise)



Note: Distribution of log MDD estimates generated by 250 repetitions of the particle learning filter applied to each model variant. For each model, distributions of log MDD estimates are generated for  $M = 100,000$  (solid, black),  $M = 10,000$  (dashed, red) and  $M = 1,000$  (dotted, blue) particles. In each case, the distributions shown are kernel density estimates obtained from the 250 simulated log MDD values. The horizontal dashed (black) line, displays the average estimate obtained with  $M = 100,000$  particles, and the surrounding shaded area demarcates bands of plus/minus twice the estimated numerical standard error, as reported in Table 5 of the paper, around that average estimate.

### R.7.4 CPI data: relative effective sample sizes

Figure R.86 reports relative effective sample sizes (ESS) for the particle learning estimates of our four model variants when estimated with 10,000 particles. Relative ESS are computed as  $\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$  where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix.


The ESS are generally quite high, often exceeding 90%. But, during the mid- and late-1970s, and to a lesser degree also during the latest recession, the relative ESS is at times very low, at least for models  $\mathcal{M}_0$ ,  $\mathcal{M}_2$ , and  $\mathcal{M}_3$ , where the gap persistence parameter  $\theta_t$  or the forecast stickiness parameter  $\lambda_t$  are time-varying. Nevertheless, with 10,000 particles, a relative ESS of one percent, i.e. 0.01, still corresponds to an ESS of 1,000 particles. As shown in Figure R.86, the relative ESS does not fall below this threshold, except for a few individual observations during the 1970s for models  $\mathcal{M}_2$  and  $\mathcal{M}_2$ , that are typically associated with particle values for  $\theta_t$  (the AR(1) gap parameter) near the unit circle.

Figure R.87 shows similar results for the relative ESS derived from our four model variants when the noise component in the inflation process is set to zero, as in (R.23). (If anything, in the absence of noise in the inflation equation, particularly low values of the relative ESS occur a little more frequently.)

### R.7.5 Forecasting CPI Inflation with Our State Space Models

This section presents results of forecast comparisons obtained from our CPI data set. Section R.1.2 describes the setup of the forecast comparison. Table R.6 presents results for models with noise in the inflation specification, table R.6 for the case without noise

Figure R.86: Relative ESS




Note: Relative effective sample sizes (ESS) with 10,000 particles for the particle learning estimates of our four model variants. Relative ESS are computed as

$$\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$$

where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix. With 10,000 particles, a relative ESS of one percent, i.e. 0.01, corresponds to an ESS of 1,000. Observations of relative ESS below one percent are marked with a diamond (red).

Figure R.87: Relative ESS (w/o noise)



Note: Relative effective sample sizes (ESS) with 10,000 particles for the particle learning estimates of our four model variants when the irregular noise component in the inflation equation has been shut off, as in (R.23). Relative ESS are computed as

$$\text{Rel ESS}_t = \frac{M}{\sum_{i=1}^M (W_t^{(i)})^2}$$

where  $W_t^{(i)}$  is the weight associated with particle  $i$  as defined in Step 3 of the particle filter described in Section III.1.2 of the supplementary appendix. With 10,000 particles, a relative ESS of one percent, i.e. 0.01, corresponds to an ESS of 1,000. Observations of relative ESS below one percent are marked with a diamond (red).

Table R.6: Inflation Forecast Comparison (CPI data)

horizon	SPF (RMSE)	RE model forecasts (rel. RMSE)				SI model forecasts (rel. RMSE)			
		$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
PANEL (A): 1981:Q2 to 2018:Q3									
1	1.15	1.45	1.03	1.40	1.44	1.47	1.09	1.42	1.48
2	1.84	1.04	1.03	1.04	1.05	1.03	1.03	1.03	1.04
3	1.92	1.00	1.01	1.01	1.01	1.00	1.00	1.00	1.00
4	1.92	0.99	0.99	1.00	1.00	1.00	1.00	1.00	1.00
5	1.99	0.99	0.98	0.99	0.99	0.99	0.99	0.99	0.99
PANEL (B): 1981:Q2 to 1984:Q4									
1	1.53	1.19	1.12	1.27	1.29	1.26	1.17	1.30	1.29
2	2.04	1.11	1.11	1.10	1.11	1.09	1.09	1.08	1.09
3	2.25	1.01	1.01	1.02	1.01	1.01	1.01	1.01	1.01
4	2.20	0.95	0.95	0.98	0.97	0.99	0.98	0.99	0.99
5	2.66	0.93	0.91	0.94	0.93	0.95	0.95	0.95	0.95
PANEL (C): 1985:Q1 to 2000:Q1									
1	0.71	1.50	1.11	1.50	1.53	1.53	1.17	1.53	1.55
2	1.24	1.04	1.03	1.03	1.05	1.04	1.02	1.03	1.03
3	1.34	0.99	0.99	0.99	1.00	0.99	0.99	0.99	0.99
4	1.39	0.99	0.99	0.99	1.00	0.99	0.99	1.00	1.00
5	1.42	0.98	0.98	0.99	0.99	0.97	0.98	0.98	0.98
PANEL (D): 2000:Q1 to 2018:Q3									
1	1.34	1.50	0.97	1.40	1.45	1.51	1.05	1.42	1.50
2	2.17	1.03	1.02	1.03	1.03	1.02	1.01	1.02	1.03
3	2.23	1.01	1.01	1.01	1.01	1.00	1.01	1.00	1.00
4	2.22	1.00	1.00	1.01	1.00	1.00	1.00	1.00	1.00
5	2.22	1.00	1.00	1.01	1.01	1.01	1.00	1.01	1.01

Note: RMSE of SPF forecasts and relative RMSE of RE and SI predictions generated by a given model compared to the SPF forecast. (Numbers below one indicate a lower RMSE of the model forecasts). In each panel, model forecasts used are based on filtered estimates using data since 1981:Q2. Forecast errors are then collected over the (sub)periods indicated in each panel.

Table R.7: Inflation Forecast Comparison (CPI data, models with noise in inflation)

horizon	SPF (RMSE)	RE model forecasts (rel. RMSE)				SI model forecasts (rel. RMSE)			
		$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$	$\mathcal{M}_0$	$\mathcal{M}_1$	$\mathcal{M}_2$	$\mathcal{M}_3$
PANEL (A): 1981:Q2 to 2018:Q3									
1	1.15	1.45	1.29	1.46	1.46	1.47	1.31	1.48	1.49
2	1.84	1.04	1.04	1.04	1.05	1.03	1.04	1.04	1.04
3	1.92	1.00	1.01	1.01	1.01	1.00	1.01	1.00	1.00
4	1.92	0.99	0.99	1.00	1.00	1.00	1.00	1.00	1.00
5	1.99	0.98	0.98	0.99	0.99	0.99	0.99	0.99	0.99
PANEL (B): 1981:Q2 to 1984:Q4									
1	1.53	1.18	1.13	1.27	1.29	1.25	1.17	1.30	1.31
2	2.04	1.10	1.11	1.09	1.11	1.08	1.09	1.08	1.08
3	2.25	1.01	1.02	1.02	1.02	1.01	1.01	1.01	1.01
4	2.20	0.93	0.94	0.97	0.96	0.99	0.98	0.98	0.99
5	2.66	0.91	0.90	0.94	0.92	0.94	0.94	0.94	0.95
PANEL (C): 1985:Q1 to 2000:Q1									
1	0.71	1.55	1.49	1.57	1.59	1.59	1.48	1.60	1.60
2	1.24	1.04	1.04	1.04	1.07	1.04	1.03	1.04	1.04
3	1.34	0.99	0.98	0.99	1.00	1.00	0.99	0.99	0.99
4	1.39	0.99	0.98	0.99	1.01	0.99	0.99	1.00	1.00
5	1.42	0.98	0.98	0.99	0.99	0.98	0.98	0.98	0.98
PANEL (D): 2000:Q1 to 2018:Q3									
1	1.34	1.48	1.28	1.48	1.47	1.50	1.31	1.50	1.51
2	2.17	1.02	1.03	1.03	1.03	1.02	1.03	1.03	1.02
3	2.23	1.00	1.02	1.01	1.01	1.00	1.01	1.00	1.00
4	2.22	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
5	2.22	1.01	1.01	1.01	1.01	1.01	1.01	1.01	1.01

Note: RMSE of SPF forecasts and relative RMSE of RE and SI predictions generated by a given model compared to the SPF forecast. (Numbers below one indicate a lower RMSE of the model forecasts). In each panel, model forecasts used are based on filtered estimates using data since 1981:Q2. Forecast errors are then collected over the (sub)periods indicated in each panel.



## R.8 Monte Carlo Study

This appendix subjects our particle learning filters to a Monte Carlo study. In this simulation exercise, many samples of artificial data are simulated and we ask two questions:

1. Does the particle learning filter for the  $\mathcal{M}_2$  model produce estimates of latent states and parameters that are unbiased and reliably close to their true values?
2. Does a comparison of log MDDs estimated for all of our four model variants, which vary in whether  $\theta$  and  $\lambda$  are assumed to be time-varying or constant, correctly detect that the data was generated by the  $\mathcal{M}_2$  model?

To keep computational cost manageable, we limited the design of the Monte Carlo study to consider only filtered estimates in answering question 1. (Smoothed estimates should prove even more reliable than what is reported here.) In addition, we limited the number of particles used by the particle learning filter to 10,000.

The Monte Carlo experiments rely on 320 independent samples of artificial data of length  $T = 200$ .<sup>16</sup> For each simulation, initial values  $\varepsilon_t$ ,  $\mathbf{F}_t\varepsilon_t$ , and  $\mathbf{F}_t\tau_t$  are drawn from their respective priors described in section 4 of the paper, and for  $\lambda_0$ , initial values are drawn from a uniform distribution with support between zero and one. Initial values for  $\theta_0$  and the log SV processes were set equal to their prior means in order to avoid ill-behaved draws, which could otherwise arise, in particular in light of the random walk processes driving the log's of the SV states. For similar reasons, the initial level of RE trend inflation  $\tau_0$  has been set equal to its prior mean.<sup>17</sup> Trajectories for the latent state variables are then simulated based on their respective laws of motions.

### R.8.1 Simulated bias in estimates

Figures R.88– R.91 depict the distribution of differences between estimated and true values obtained from the particle learning filter of  $\mathcal{M}_2$  applied to each simulated sample of data. Overall, estimation errors are centered around zero and reasonably closely so, reflecting the filtered nature of the estimates and the relatively short sample size.<sup>18</sup>


<sup>16</sup>Our original data set has 201 observations of U.S. data.

<sup>17</sup>The RE trend,  $\tau_t$ , constitutes the Beveridge-Nelson trend of  $\mathbf{F}_t\tau_t$ . Differences between RE and SI trend evolve as a stable AR(1) as described by equation 18 in the paper, so that differences in initial conditions of  $\tau_0$  and  $\mathbf{F}_0\tau_0$  dissipate. Drawing  $\mathbf{F}_0\tau_0$  from its prior does not generate ill-behaved draws. In fact, setting all initial values of the linear states equal to their prior means would imply an identity of RE and SI components of the linear states at  $t=0$  suggesting that  $\lambda_0 = 0$  (contrary to the actual level of  $\lambda_0$  drawn to initialize the simulated trajectory of  $\lambda_t$ ). In an earlier version of this paper, simulated initial levels for trend and gap inflation had been fixed at their prior means, leading to a seeming bias in estimates of  $\lambda_t$  early in the sample, which turns out to have been due to the aforementioned set up of the earlier simulation design. We would like to thank an anonymous reviewer for the encouragement to investigate this issue further.


<sup>18</sup>The only exception is some initial bias in SV levels, shown in figure R.89, which we attribute to Jensen inequality effects in transforming the priors for log variances to SV levels. We also simulated the bias of *smoothed* estimates, albeit with a smaller number of simulations in light of the increased computational cost; for smoothed SV estimates, we found the bias in initial SV estimates to dissipate.

Figure R.88: Simulated Bias in Particle Learning Estimates of Trend and Gap Inflation

(a) Trend inflation

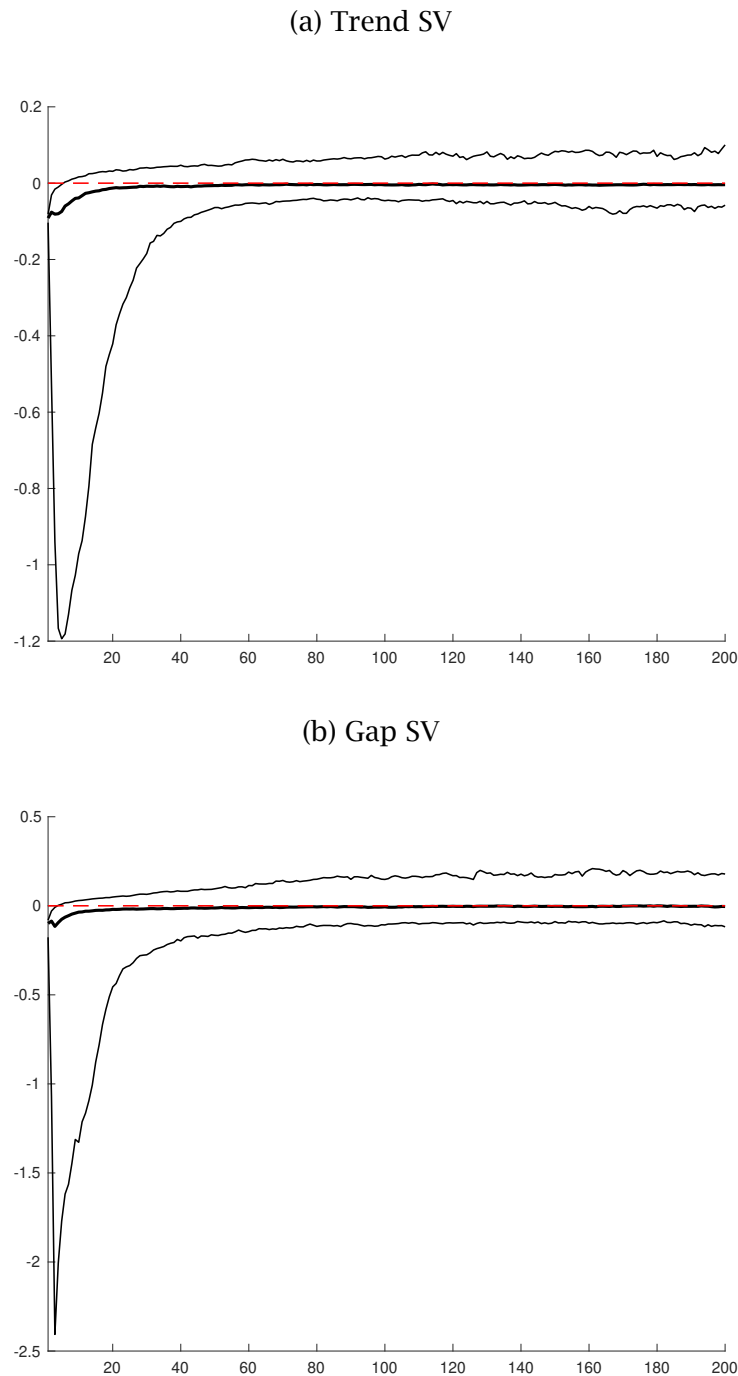


(b) Gap inflation



Note: Distribution of differences between particle learning estimates and true values of trend inflation,  $\tau_t - \tau_{t|t}$ , and gap inflation,  $\varepsilon_t - \varepsilon_{t|t}$ , respectively. The distribution of differences is generated over 640 samples of simulated data of length  $T = 200$  from the  $\mathcal{M}_2$  model and the corresponding particle learning filter using 10,000 particles. Thick solid lines denote the median, and thin lines the 5% and 95% quantiles of the simulated distribution.


Figure R.89: Simulated Bias in Particle Learning Estimates of Stochastic Volatilities




Note: Distribution of differences between particle learning estimates of stochastic volatilities and true values,  $\zeta_{i,t} - \zeta_{i,t|t}$  for  $i = \eta, \nu$ . The distribution of differences is generated over 640 samples of simulated data of length  $T = 200$  from the  $\mathcal{M}_2$  model and the corresponding particle learning filter using 10,000 particles. Thick solid lines denote the median, and thin lines the 5% and 95% quantiles of the simulated distribution.

Figure R.90: Simulated Bias in Particle Learning Estimates  $\theta_t$  and  $\lambda_t$

(a) AR(1) coefficient  $\theta_t$





(b) SI parameter  $\lambda_t$



Note: Distribution of differences between particle learning estimates and true values of  $\theta_t$ ,  $\theta_t - \theta_{t|t}$ , and  $\lambda_t$ ,  $\lambda_t - \lambda_{t|t}$ , respectively. The distribution of differences is generated over 640 samples of simulated data of length  $T = 200$  from the  $\mathcal{M}_2$  model and the corresponding particle learning filter using 10,000 particles. Thick solid lines denote the median, and thin lines the 5% and 95% quantiles of the simulated distribution.



Figure R.91: Simulated Bias in Particle Learning Estimates of Static Parameters

(a) Variance of Shocks to log Trend SV,  $\sigma_\eta^2$  (b) Variance of Shocks to log Trend SV,  $\sigma_v^2$



(c) Variance of Shocks to  $\theta_t$ ,  $\sigma_\phi^2$

(d) Variance of Shocks to  $\lambda_t$ ,  $\sigma_\kappa^2$



Note: Distribution of differences between particle learning estimates and true values of static parameters. The distribution of differences is generated over 640 samples of simulated data of length  $T = 200$  from the  $\mathcal{M}_2$  model and the corresponding particle learning filter using 10,000 particles. Thick solid lines denote the median, and thin lines the 5% and 95% quantiles of the simulated distribution.

## R.8.2 Model detection

Turning to the question of model selection, Figures R.92 and R.93 report detection rates for different models obtained in simulated data generated by  $\mathcal{M}_2$ . This setup assumes that the data generating process features a time-varying persistence parameter,  $\theta_t$ , for gap inflation, and a time-varying SI parameter  $\lambda_t$ . For a given pair of models, we measure the share of simulations where the difference in log MDDs was larger than two or smaller than minus two. Log MDDs are obtained by applying particle learning filters for  $\mathcal{M}_0$ ,  $\mathcal{M}_1$ ,  $\mathcal{M}_2$ ,  $\mathcal{M}_3$  to the simulated data. As before, 640 samples of data of length  $T = 200$  were simulated from the  $\mathcal{M}_2$  model. For each sample, particle learning filters for each model are run with 10,000 particles. In each case, the particle learning filter is run ten times and the average log MDD across these ten repetitions is used


Figure R.92 presents detection obtained from comparing the true model,  $\mathcal{M}_2$ , where  $\lambda_t$  and  $\theta_t$  are time-varying, against the three alternative models. Given the relatively short sample size of  $T = 200$ , the ability to distinguish between  $\mathcal{M}_2$  and the alternative  $\mathcal{M}_0$  and  $\mathcal{M}_3$ , where time-variation is switched off in only one of the two parameters  $\theta_t$  and  $\lambda_t$ , is much more limited. Overall, in the presence of joint time-variation in  $\theta_t$  and  $\lambda_t$  in samples of typical length, the log MDDs generated by the particle learning filter seem most capable to detect failures of the  $\mathcal{M}_1$  model, which assumes both parameters to be constant, while the true  $\mathcal{M}_2$  model is much harder to distinguish from the other two alternatives,  $\mathcal{M}_1$  and  $\mathcal{M}_3$ , that treat at least one of the two parameters in question as constant. These results caution against the ability to strongly distinguish between all model alternatives in typical macroeconomic data. However, they also suggest that the strength of the log MDDs in favor of  $\mathcal{M}_2$  generated from the actual data, as reported in the paper, is quite remarkable.

The simulated detection rates reported thus far were obtained from environments that assume a sizable volatility of noise in inflation and measurement errors in survey expectations.<sup>19</sup> Of course, measurement error also gets in the way of detecting nuanced model differences, such as the extent of time-variation in inflation gap persistence and forecast stickiness. As an alternative cross-check on the performance of our estimation setup, we have also computed detection rates obtained from simulated data sets, where the degree of measurement error is much smaller — specifically, volatility of the measurement error in every observation equation is reduced to 0.01 as opposed to the baseline value of  $\sqrt{0.1} \approx 0.32$ .<sup>20</sup> As shown in Figures R.94, detection rates in favor of the true model,  $\mathcal{M}_2$ , are considerably larger in this environment. As before, the model variant where both  $\theta$  and  $\lambda$  are assumed to be constant,  $\mathcal{M}_1$ , is soundly rejected in most of the simulated data samples when compared to any of the other alternatives, where at least one of the two parameters is treated as time-varying.

<sup>19</sup>Henceforth, for brevity, we refer to the noise component in inflation and measurement error in the survey equations as measurement error. All simulations use a measurement error volatility equal to the prior mean of  $\sqrt{0.1}$  used in each estimation.


<sup>20</sup>That means the volatility of measurement errors in inflation and surveys, as in equations (1.1) and (3.1) of the paper, is set to 0.01. The priors placed on these variance parameters are re-centered around the true value as well.

Figure R.92: Detection Rates of True Model




Note: Solid (red) lines denote the share of simulations (in percentage points) where the difference in log MDDs between the pair of models considered was larger than two. Dashed (black) lines report the corresponding share of simulations where the difference in log MDD was smaller than minus two. The log MDDs are generated over 640 samples of simulated data of length  $T = 200$  from the  $M_2$  model. For each sample, particle learning filters with 10,000 particles are run estimating models  $M_0, M_1, M_2, M_3$ . In each case, the particle learning filter is run ten times and the average log MDD across these ten repetitions is used. The x-axis measures time  $t = 1, \dots, 200$ .

Figure R.93: Detection Rates against  $M_1$  Model




Note: Solid (red) lines denote the share of simulations (in percentage points) where the difference in log MDDs between the pair of models considered was larger than two. Dashed (black) lines report the corresponding share of simulations where the difference in log MDD was smaller than minus two. The log MDDs are generated over 640 samples of simulated data of length  $T = 200$  from the  $M_2$  model. For each sample, particle learning filters with 10,000 particles are run estimating models  $M_0, M_1, M_2, M_3$ . In each case, the particle learning filter is run ten times and the average log MDD across these ten repetitions is used. The x-axis measures time  $t = 40, \dots, 200$ . (The first 39 observations are omitted from each figure, since detection rates remain near zero over the equivalent of the first ten years of quarterly observations in every case.)

Figure R.94: Detection Rates of True Model when Measurement Error is Small



Note: Solid (red) lines denote the share of simulations (in percentage points) where the difference in log MDDs between the pair of models considered was larger than two. Dashed (black) lines report the corresponding share of simulations where the difference in log MDD was smaller than minus two. The log MDDs are generated over 640 samples of simulated data of length  $T = 200$  from the  $\mathcal{M}_2$  model. For each sample, particle learning filters with 10,000 particles are run estimating models  $\mathcal{M}_0, \mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3$ . In each case, the particle learning filter is run ten times and the average log MDD across these ten repetitions is used. The x-axis measures time  $t = 40, \dots, 200$ . (The first 39 observations are omitted from each figure, since detection rates remain near zero over the equivalent of the first ten years of quarterly observations in every case.)

Figure R.95: Detection Rates against  $\mathcal{M}_1$  Model when Measurement Error is Small



Note: Solid (red) lines denote the share of simulations (in percentage points) where the difference in log MDDs between the pair of models considered was larger than two. Dashed (black) lines report the corresponding share of simulations where the difference in log MDD was smaller than minus two. The log MDDs are generated over 640 samples of simulated data of length  $T = 200$  from the  $\mathcal{M}_2$  model. For each sample, particle learning filters with 10,000 particles are run estimating models  $\mathcal{M}_0, \mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3$ . In each case, the particle learning filter is run ten times and the average log MDD across these ten repetitions is used. The x-axis measures time  $t = 1, \dots, 200$ .



## References

- Cogley, T., Primiceri, G. E., and T. J. Sargent (2010), 'Inflation-gap persistence in the U.S.', *American Economic Journal: Macroeconomics* 2(1), 43–69.
- Cogley, T. and T.J. Sargent (2015), "Measuring price-level uncertainty and instability in the U.S., 1850–2012." *Review of Economics and Statistics*, 97, 827–838.
- Creal, D. (2012). A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31, 245–296.
- Del Negro, M., Hasegawa, R. B., and F. Schorfheide (2016), 'Dynamic prediction pools: An investigation of financial frictions and forecasting performance', *Journal of Econometrics* 192(2), 391–405.
- Faust, J. and J.H. Wright (2013), "Forecasting inflation." In *Handbook of Economic Forecasting*, vol. 2, Ch. 1 (Elliot, G. and A. Timmermann, eds.), 2–56, Elsevier Science.
- Fuentes-Albero, C. and L. Melosi (2013), 'Methods for computing marginal data densities from the Gibbs output', *Journal of Econometrics* 175(2), 132–141.
- Geweke, J. (1989), 'Bayesian inference in econometric models using Monte Carlo integration', *Econometrica* 57(6), 1317–1339.
- Grassi, S. and T. Proietti (2010), "Has the volatility of U.S. inflation changed and how?" *Journal of Time Series Econometrics*, 2:1, article 6.
- Hamilton, J. D. (1994), *Time-Series Analysis*, Princeton University Press, Princeton, NJ.
- Hansen, L. P. and T.J. Sargent (2007), *Robustness*, Princeton University Press.
- Herbst, E. and F. Schorfheide (2014), 'Sequential Monte Carlo sampling for DSGE models', *Journal of Applied Econometrics* 29(7), 1073–1098.
- Koop, G. (2003), *Bayesian Econometrics*, Wiley-Interscience.
- Mertens, E. (2016), "Measuring the level and uncertainty of trend inflation." *The Review of Economics and Statistics*, 98(5), 950–967.
- Shephard, N. (2013), "Martingale unobserved component models." Economics Series Working Papers 644, Department of Economics, University of Oxford.
- Stock, J.H. and M.W. Watson (2007), "Why has US inflation become harder to forecast?" *Journal of Money, Credit and Banking*, 39(S1), 3–33.
- Stock, J.H. and M.W. Watson (2010), "Modeling inflation after the crisis." In *MACROECONOMIC CHALLENGES: THE DECADE AHEAD*, Chapter 3, (C.S. Hakkio and E.S. Knotek, II, eds.), 173–220, Federal Reserve Bank of Kansas City.