# Particle Filtering, Learning, and Smoothing for Mixed-Frequency State-Space Models

December 3, 2018

**Abstract**

We propose a flexible smoothing-based particle filter for general mixed-frequency state-space models. Our approach employs a backward smoother to filter high-frequency state variables from low-frequency observations. Moreover, it preserves the sequential nature of particle filters, allows for non-Gaussian shocks and nonlinear state-measurement relation, and alleviates the concern of sample degeneracy. In the empirical study, it is used to estimate predictive regressions for Treasury bond and US dollar index returns with quarterly predictors and monthly stochastic volatility. Empirically, stochastic volatility improves model inference and forecasting power in a mixed-frequency setup but not in quarterly aggregate models.

JEL Classification: C13, C32, C53.

Keywords: Mixed-Frequency, State-Space Models, Particle Filtering, Particle Learning, Stochastic Volatility, Return Predictability

# 1 Introduction

In economics and finance, state-space models have become a popular modeling device over the last two decades. When estimating these models, we face the tightly coupled problems of state and parameter inference. Often we are interested in sequential real-time estimation, where estimates of state variables and forecasts are formed as new information arrives. However, at least two dilemmas arise in state-space modeling. First, economic data is often sampled at different frequencies. For instance, GDP growth rates and many surveys of macroeconomic variables are available only at a quarterly frequency, whereas other variables, including asset returns and interest rates, are available at much higher frequencies. Second, even for the same variable, the sampling frequency may differ across periods. For instance, the real consumption expenditure per capita is collected annually before 1947, quarterly until 1959, and monthly afterward. Simple treatments of such unbalanced datasets either adopt a time aggregation to match a common low sampling frequency, or they exclusively focus on a uniform subset of the entire dataset, or even use a rather crude data augmentation procedure. However, such simplifications may destroy useful information that we could potentially gain from high-frequency observations. Well-known high-frequency empirical features include, for instance, time-varying volatility. Consequently, ignoring this information may give rise to inconsistency and biased model inference, and incur economic losses for investors and policymakers who make decisions based on these data.

In light of these drawbacks, there is a rising interest in econometric methods designed to handle mixed-frequency data. We contribute to this growing body of literature by proposing a simple yet general particle filtering framework for mixed-frequency state-space models (MFSSMs). The central issue that lies in the estimation of MFSSMs is the filtering of state variables. Among all filtering techniques, the Kalman filter, due to its analytical tractability, is the most commonly used approach; see, for instance, Harvey (1989), Mariano and Murasawa (2003), Giannone et al. (2008), Aruoba et al. (2009), and Durbin and Koopman (2012). However, the Kalman filter builds on the assumptions of normality and linearity of the state-measurement relation. These assumptions are too restrictive given the empirically observed skewness and kurtosis of economic data, as well as the time variation in their volatility. Efficient alternatives include the Markov-Chain Monte Carlo (MCMC) method, which has been used by Schorfheide and Song (2015) and Marcellino et al. (2016), but are not tailor-made for sequential real-time inference. In contrast, particle filters alleviate these concerns. When the filtering density is not available in closed form, the sequential inference is achieved by using Monte-Carlo samples to approximate filtering distributions. Without loss of generality and given its well-known advantage of improving the filtering efficiency, we build the mixed-frequency particle filter (MFPF) on the resample-propagate scheme of Carvalho et al. (2010).

The problem of mixed sampling frequencies poses additional difficulties in sequential filtering. Simple implementations employ a state space augmentation step to convert MFSSMs to synchronous-evolving models; see, for instance, Schorfheide et al. (2018). However, they are potentially hampered by sample degeneracy, which refers to the notorious fact that state variables are improperly represented by only a small effective number of Monte-Carlo samples; see Doucet et al. (2000) for a discussion.

We take a novel view toward MFSSMs that mitigates sample degeneracy. We realize that, when we cast low-frequency observations in a state-space form, they tend to be jointly determined by lagged state variables with the length equal to the sampling interval. Proceeding with the forward filter requires the joint smoothing distribution of lagged state variables. We thus employ a smoothing-based technique to proceed when low-frequency observations become available. In our implementation, we apply the backward smoother, which has been examined by Frühwirth-Schnatter (1994) and Godsill et al. (2004). The backward smoother minimizes the risk of sample degeneracy due to mixed-frequency data by rejuvenating Monte-Carlo samples that represent lagged state variables. Further, the backward smoother preserves the sequential nature of particle filters when low-frequency observations are incorporated. As the notion of smoothing-based filtering is not restricted to a specific filter or smoother, it can be immediately extended to more general setups of particle filtering. We also discuss extensions that allow for sequential parameter learning.

To illustrate the empirical advantage of mixed-frequency models and MFPF, we use the Survey of Professional Forecasters (SPF) to predict asset returns. SPF is provided by the Philadelphia Fed and aggregates professional forecasts of key economic variables. Recent literature, including Chun (2010), Chernov and Mueller (2012), Eriksen (2017), and Foroni et al. (2018), shows that survey data predict key financial asset returns. These survey variables are only updated at a quarterly frequency. In predicting asset returns, the usual practice is to temporally aggregate them to a quarterly frequency. However, asset returns exhibit time-varying volatility. This time-variation becomes weaker at lower frequencies such as quarterly. However, ignoring time-varying volatility might give biased model inference and incur economic losses for investors who manage portfolios and volatility at a relatively high frequency. Thus, there is a practical demand for incorporating quarterly survey variables in monthly-evolving predictive regressions that possibly preserve the high-frequency nature of time-varying volatility.

We use the quarterly survey of the growth rates of industrial production, real consumption expenditure per capita, and CPI inflation rates to predict monthly returns of the Treasury bonds and US trade-weighted dollar index. Incorporating time-varying volatility results in state-space nonlinearity and prohibits the use of the Kalman filter. For model estimation, we embed MFPF to a random-walk Metropolis-Hasting algorithm, where we use MFPF for filtering and likelihood computation and an MCMC iterator for generating parameter posteriors. Empirically, we find that, at a monthly frequency, incorporating stochastic volatility improves return forecasts by the quarterly survey variables, with more favorable density forecasts, Akaike information criteria, and prediction $R^2$-values. This finding justifies the usefulness of employing mixed-frequency models to forecast monthly returns using quarterly predictors. Further, we also examine quarterly aggregate models, in which quarterly predictors are used to forecasting quarterly returns. We find that, at a quarterly frequency, stochastic volatility does not improve model inference or return forecast. Economically, this implies that investors cannot take advantage of volatility timing at a quarterly frequency.

This paper contributes to the literature of Bayesian mixed-frequency approaches along various dimensions. Most of the literature considers modeling mixed-frequency data using linear-Gaussian models. Applications include, e.g., Mariano and Murasawa (2003) who construct a monthly GDP series from quarterly GDP and business cycle variables, and Aruoba et al. (2009) who develop an economic activity index in real time from various mixed-frequency series. Giannone et al. (2008) evaluate the marginal impact of monthly data releases on nowcasts of quarterly real GDP growth rates. Estimation of these models enjoys closed-form solutions known as Kalman filter; see Harvey (1989) and Durbin and Koopman (2012) for a textbook treatment. Marcellino et al. (2016) develop a mixed-frequency GDP forecasting model with stochastic volatility and employ an MCMC approach for estimation. Schorfheide and Song (2015) develop a Bayesian mixed-frequency VAR, coupled with Minnesota prior and MCMC estimators, to forecast quarterly macroeconomic variables. This paper proposes a flexible sequential particle filtering framework for MFSSMs that nests all these model specifications and further allows non-Gaussian and nonlinear dynamics.

We take a novel view that mitigates sample degeneracy due to mixed-frequency data. Among the literature dealing with mixed sampling frequencies, Mariano and Murasawa (2003) modify the measurement equation by setting the loadings on state variables to zero. Giannone et al. (2008) set the measurement error of low-frequency observations to infinity. Aruoba et al. (2009) and Schorfheide and Song (2015) vary the dimension of the observations. In identifying long-run risks from mixed-frequency consumption data and in predicting stock returns, Schorfheide et al. (2018) and Leippold and Yang (2018) employ a state space augmentation procedure to convert their models to synchronous-evolving models, respectively. We employ backward smoothers to deal with mixed-frequency data, given that the lagged stable variables jointly determine low-frequency observations. The smoothing-based approach, as widely recognized in the Bayesian literature, alleviates the notorious concern of sample degeneracy. Another branch of literature dealing with mixed-frequency data takes an observation-driven approach; see the seminal work of Ghysels et al. (2004) and Ghysels et al. (2007) for MIDAS. Bai et al. (2013) examine the relation between MIDAS and state-space models. Ghysels (2016) generalizes MIDAS to a VAR setup. Further, it allows the prediction of both high- and low-frequency observations from mixed-frequency data. Valle e Azevedo et al. (2006) propose a factor model that allows for trend-cycle decomposition and mixed-frequency macroeconomic data. Creal et al. (2014) develop an observation-driven mixed-frequency factor model.

# 2 Mixed-Frequency State-Space Models

In what follows, we introduce our state-space model framework that allows for mixed-frequency variables, and we present a detailed discussion of its connection to our benchmark, the synchronously evolving state-space models.

## 2.1 Model Setup

The standard state-space model takes the form

$$\begin{aligned} X_{t+1} &\sim p\left(X_{t+1} \middle| X_t, \Theta\right), \\ Y^H_{t+1} &\sim p\left(Y^H_{t+1} \middle| X_{t+1}, \Theta\right), \end{aligned} \tag{1}$$

where $X$ is the set of state variables driving the stochastic dynamics, $Y^H$ is the stream of observations available per unit of time, and $\Theta$ is the set of parameters and fixed as constant. In particular, $H$ is used as the superscript to emphasize that $Y^H$ is observed at the basis frequency, namely, the frequency at which the state variable evolves. Literally, $H$ denotes a relatively high frequency, which we will distinguish from a lower frequency denoted by $L$.

In the macro-finance literature, the basis frequency is often one month, at which various economic variables including industrial production growth rates, inflation rates, asset returns, and others, are available. There are also variables measured only at a quarterly frequency, for instance, GDP growth rates and data from the Survey of Professional Forecasters. The Survey of Professional Forecasters is conducted by the Philadelphia Fed and summarizes the aggregate forecasts of the ongoing and future economy, which we will use to illustrate the empirical performance of mixed-frequency predictive regressions for asset returns.

Researchers often want to model the joint dynamics of both low- and high-frequency variables. The combination of observations available at different frequencies is motivated by the fact that low(high)-frequency variables can often improve the forecast of high(low)-frequency variables. We denote the low-frequency observation by $Y^L$, which is assumed to be available per $M$ units of time only. Bayesian approaches involve incorporating $Y^L$ in a way that preserves the high-frequency dynamics in equation (1). Often, these approaches imply that $Y^L$ is jointly determined by the entire path of the state variable over the recent $M$ periods

$$Y^L_{t+M} \sim p\left(Y^L_{t+M} \middle| X_{t+1:t+M}, \Theta\right), \tag{2}$$

with $X_{t+1:t+M} = (X_{t+1}, ..., X_{t+M})$. We call the joint equations (1) and (2) mixed-frequency state-space models (MFSSMs). We use $Y_t$ to denote the set of all observations at time $t$, and $Y^t$ the set of all observations available up to time $t$.

Noteworthy, equations (1) and (2) place minimum model restrictions and nest various Bayesian mixed-frequency models examined by the literature; see, for instance, Mariano and Murasawa (2003), Aruoba et al. (2009), Schorfheide and Song (2015), Marcellino et al. (2016), Schorfheide et al. (2018), and Leippold and Yang (2018), to name a few. Therefore, in presenting our approaches, we do not impose any parametric structure. With slight modifications, the MFSSM framework can also handle more general situations such as time-varying sampling frequencies or temporally missing data, but this paper focuses on the generic form given in equations (1) and (2) to illustrate the notion of smoothing-based particle filtering.

A well-established literature addresses the issue of mixed sampling frequencies. Among them, the Kalman filter is the most commonly used Bayesian approach. However, the model specifications for which the Kalman filter is applicable are too restrictive. First, it assumes the random shocks in the state-space models to be Gaussian. However, this assumption is most often violated in economics, as skewed distributions and heavy tails regularly characterize the data. Second, it imposes a linear state-measurement dependence structure that prohibits one from incorporating more flexible yet realistic model features such as lognormally distributed stochastic variance. In contrast, particle filters allow all these features by using Monte-Carlo samples as proxies for the filtering distributions that are not available in closed form. Efficient alternatives include the Markov Chain Monte Carlo (MCMC) method (see, for instance, Marcellino et al. (2016)), but are often not tailor-made for sequential real-time inference. In what follows, we present a way to deal with mixed-frequency data in a particle filter framework.

# 3 Particle Filtering

Estimating MFSSMs hinges on the filtering of state variables when observations are available at different frequencies. In this regard, the literature takes diverging views. For example, Mariano and Murasawa (2003) modify the measurement equations by setting the loadings on the state variables to zero. Giannone et al. (2008) set the measurement error of low-frequency observations to infinity. Aruoba et al. (2009), Durbin and Koopman (2012), and Schorfheide and Song (2015) vary the vector of observations in different periods. Schorfheide et al. (2018) and Leippold and Yang (2018) adopt a state space augmentation approach. To implement particle filters for MFSSMs in an efficient way, we adopt a novel view. In light of the dependence structure of low-frequency variables on lagged state variables, as illustrated by equation (2), we employ smoothing techniques to deal with mixed-frequency data. The use of smoothing techniques greatly mitigates the notorious concern of sample degeneracy that may haunt accurate filtering of MFSSMs. Moreover, it preserves the sequential nature of particle filters and is convenient for sequential inference. We focus on particle filtering of MFSSMs, but also discuss the use of smoothing techniques in other scenarios with mixed-frequency data, for example, sequential parameter learning.

## 3.1 Resample-Propagate Approach

To fix the idea, we restrict our analysis to the resample-propagate framework proposed by Carvalho et al. (2010). As smoothing is largely independent of the context in which the exact particle filters are employed, the smoothing-based approach for MFSSMs can be extended immediately to alternative filters, including those proposed by Gordon et al. (1993), Liu and West (2001), and Storvik (2002), to name a few.

Filtering of MFSSMs is achieved by extending the standard resample-propagate filter. For synchronous-evolving models like in equation (1), the resample-propagate procedure obtains the time-$(t+1)$ filtering distribution from that at time $t$ based on the Bayes' rule

$$\begin{aligned} p\left(X_{t+1} \middle| Y^{t+1}, \Theta\right) &\propto \int p\left(Y_{t+1} \middle| X_t, \Theta\right) p\left(X_{t+1} \middle| X_t, Y_{t+1}, \Theta\right) \\ &\times \; dp\left(X_t \middle| Y^t, \Theta\right), \end{aligned} \tag{3}$$

where $p\left(Y_{t+1} \middle| X_t, \Theta\right)$ is the predictive likelihood and $p(X_{t+1}|X_t, Y_{t+1}, \Theta)$ is the density of $X_{t+1}$ conditional on $X_t$ and $Y_{t+1}$. Equation (3) is fully adapted as the filtering distribution at time $t+1$ can be obtained directly from that in the last period. Monte-Carlo samples are used to approximate the state densities used in equation (3) if they are not available in closed form. Often, this is true in the presence of non-Gaussian shocks or nonlinearity of the state-space models. To implement the resample-propagate approach, we resample random draws from $p\left(X_t \middle| Y^t, \Theta\right)$ with weights proportional to the predictive likelihood, and we obtain $p\left(X_{t+1} \middle| Y^{t+1}, \Theta\right)$ by sampling from the conditional state density associated with each resampled draw. To filter the entire path of state variables, we repeat this procedure until the end of the sample period.

The resample-propagate procedure is a slight modification of the auxiliary particle filter (APF) of Pitt and Shephard (1999). The APF uses an importance function $p\left(Y_{t+1} | \alpha_{t+1} = g(x_t)\right)$ in the resampling procedure based on a best guess of $X_{t+1}$ defined by $\alpha_{t+1} = g(x_t)$, whereas the resample-propagate procedure relies on the predictive likelihood and the conditional state density. If they are available for evaluation and sampling, the efficiency of particle filter can be greatly improved. The reason is that we can use the new observation $Y_{t+1}$ in and propagation, thereby making the random draws of the state variables directed toward the observations as these draws evolve.

## 3.2 Mixed-Frequency Particle Filtering

To illustrate the smoothing-based approach to dealing with mixed sampling frequencies, we focus on the time interval $[t, t+M]$, where the low-frequency variable $Y^L$ is only observed at time $t+M$. Prior to time $t+M$, the situation is exactly the same as the synchronous-evolving state-space models, as only $Y^H$ is observed. Therefore, for $l = 0, ..., M-2$ and starting from the time-$t$ filtering distribution, we can immediately obtain, by adapting equation (3) to MFSSMs, the filtering distribution in the next period

$$
\begin{aligned}
p\left(X_{t+l+1} \middle| Y^{t+l+1}, \Theta\right) \;\; \propto & \;\; \int p\left(Y_{t+l+1} \middle| X_{t+l}, \Theta\right) p\left(X_{t+l+1} \middle| X_{t+l}, Y_{t+l+1}, \Theta\right) \\
\times & \;\; dp\left(X_{t+l} \middle| Y^{t+l}, \Theta\right). \qquad (4)
\end{aligned}
$$

However, when proceeding from time $t+M-1$ to $t+M$, equation (4) often cannot be implemented as the conditional likelihood of $Y_{t+M}$ and the conditional density of $X_{t+M}$ do not only depend on the previous state $X_{t+M-1}$, but also on its entire path over the last $M-1$ periods, i.e., on $X_{t+1:t+M-1}$. Fortunately, in most of the applications examined by the literature, the filtering algorithm can still proceed if the state space is augmented by lagged state variables. The filtering equation then becomes

$$\begin{aligned} p\left(X_{t+M}\middle|Y^{t+M},\Theta\right) &\propto \int p\left(Y_{t+M}\middle|X_{t+1:t+M-1},\Theta\right) p(X_{t+M}|X_{t+1:t+M-1},Y_{t+M},\Theta) \\ &\times \; dp(X_{t+1:t+M-1}|Y^{t+M-1},\Theta). \end{aligned} \tag{5}$$

Casting the state-augmented models into particle filters is practically equivalent to the fixed-lag smoother of Kitagawa (1996), employed by Schorfheide et al. (2018) to estimate long-run risks models and by Leippold and Yang (2018) to estimate predictive regressions for stock returns. To obtain the smoothing distribution, draws of lagged state variables are resampled in conjunction with the forward filter. However, this simplified algorithm may suffer severely from the problem of sample degeneracy, particularly when $M$ is large. The reason is that, as the simulation evolves, draws of lagged state variables are resampled sequentially in time without rejuvenation, thereby making the smoothing state distribution represented potentially improperly by only a small number of heterogeneous draws; see Doucet et al. (2000) for a discussion. In light of this drawback, we are motivated to employ more efficient smoothers to obtain the smoothing distribution $p\left(X_{t+1:t+M-1}\middle|Y^{t+M-1},\Theta\right)$.

## 3.3 Backward Smoothing

We show that in attempting to obtain $p\left(X_{t+1:t+M-1}\middle|Y^{t+M-1},\Theta\right)$, the backward smoother of Frühwirth-Schnatter (1994) and Godsill et al. (2004) serves as a convenient alternative to the fixed-lag smoother and preserve the sequential nature of particle filters. The backward smoother uses the Bayes' rule and the Markovian structure of the state process to obtain the following backward recursive representation:

$$\begin{aligned} p\left(X_{t+1:t+M-1}\middle|Y^{t+M-1},\Theta\right) &= p\left(X_{t+M-1}\middle|Y^{t+M-1},\Theta\right) \\ &\times \prod_{l=1}^{M-2} p\left(X_{t+l}\middle|X_{t+l+1},Y^{t+l},\Theta\right). \end{aligned} \tag{6}$$

To convert equation (6) into a numerically implementable form, we exploit Bayes' rule once more and express the conditional state density on the right-hand side as

$$p\left(X_{t+l}\middle|X_{t+(l+1)},Y^{t+l},\Theta\right) \propto p\left(X_{t+(l+1)}\middle|X_{t+l},\Theta\right) p\left(X_{t+l}\middle|Y^{t+l},\Theta\right). \tag{7}$$

Thus, conditional on $X_{t+(l+1)}$ and $Y^{t+l}$, the state density in the previous period is proportional to the state transition density times the filtering density, which is directly available from MFPF. In the context of particle filtering, given each draw of $X_{t+(l+1)}$, the backward smoother can be performed by recursively resampling from the filtering distribution $p\left(X_{t+l}\middle|Y^{t+l},\Theta\right)$, with sampling weights proportional to $p\left(X_{t+(l+1)}\middle|X_{t+l},\Theta\right)$. Repeating this exercise backward until time $t+1$ and piecing all resampled draws together give the smoothing density $p(X_{t+1:t+M-1}|Y^{t+M-1},\Theta)$ from $p(X_{t+M-1}|Y^{t+M-1},\Theta)$. The advantage is that random draws representing the lagged states are rejuvenated, in contrast to the fixed-lag smoother that does only resampling. Hence, the heterogeneity of Monte-Carlo samples approximating the smoothing distribution is greatly improved, as widely recognized by the Bayesian literature; see, for instance, Doucet et al. (2000) and Godsill et al. (2004). The backward smoother imposes minimum model assumptions and can handle all applications for which the state transition density is available. Moreover, the backward smoother does not build on a specific filter and can thus be fit to other filters.

### 3.4 Implementation

We denote the set of $N$ random draws from the filtering distribution $p\left(X_t \middle| Y^t, \Theta\right)$ by $\{X_t^{(i)}\}_{i\in\mathcal{N}}$, where $\mathcal{N} = \{1, 2, ..., N\}$ is the index set. Further, for notational convenience, we denote the empirical distribution formed by these draws by $p^N\left(X_t \middle| Y^t, \Theta\right)$.

**Step 1 (Filter $X_{t+l+1}$ sequentially for $l = 0, ..., M-2$).** This step executes equation (4) numerically.

**Resample.** Draw a size-$N$ set of index $(k^{(i)})_{i\in\mathcal{N}}$ from $\mathcal{N}$, with weights proportional to the predictive likelihood

$$w_{t+l+1|t+l}^{(i)} \propto p\left(Y_{t+l+1} \middle| X_{t+l}^{(i)}, \Theta\right), \tag{8}$$

where $Y_{t+l+1} = Y_{t+l+1}^H$ and $k^{(i)}$ is the index from the $i$-th draw.

**Propagate.** For each draw $k^{(i)}$, draw $X_{t+l+1}^{(i)}$ from the conditional density

$$X_{t+l+1}^{(i)} \sim p\left(X_{t+l+1} \middle| X_{t+l}^{(k^{(i)})}, Y_{t+l+1}, \Theta\right). \tag{9}$$

The set $\left\{X_{t+l+1}^{(i)}\right\}_{i\in\mathcal{N}}$ represents $N$ Monte-Carlo draws from $p\left(X_{t+l+1} \middle| Y_{t+l+1}, \Theta\right)$ and is denoted by $p^N\left(X_{t+l+1} \middle| Y_{t+l+1}, \Theta\right)$.

**Step 2 (Smooth $X_{t+1:t+M-1}$).** This step executes the backward smoother in equations (6)-(7). Set $\widetilde{X}_{t+M-1}^{(i)} = X_{t+M-1}^{(i)}$ to initialize. Smoothing is achieved by the following resampling scheme backward in time for $l = M-2, M-3, ..., 1$.

**Resample.** For each $i \in \mathcal{N}$, draw a sample, denoted by $\widetilde{X}_{t+l}^{(i)}$, from the empirical time-$(t+l)$ filtering distribution $\left\{X_{t+l}^{(j)}\right\}_{j\in\mathcal{N}}$ with weights defined as follows

$$w_{t+l|t+l+1}^{(j)} \propto p\left(\widetilde{X}_{t+l+1}^{(i)} \middle| X_{t+l}^{(j)}, \Theta\right). \tag{10}$$

The resampling procedure gives $\left\{\widetilde{X}_{t+l}^{(i)}\right\}_{i\in\mathcal{N}}$. After reaching time $t+1$ through backward smoothing, we piece these samples together to obtain

$$\left\{\widetilde{X}_{t+1:t+M-1}^{(i)}\right\}_{i\in\mathcal{N}} = \left\{\left(\widetilde{X}_{t+1}^{(i)}, ..., \widetilde{X}_{t+M-1}^{(i)}\right)\right\}_{i\in\mathcal{N}},$$

which is also the empirical smoothing distribution $p^N(X_{t+1:t+M-1}|Y^{t+M-1}, \Theta)$.

**Step 3 (Filter $X_{t+M}$).** Implementing the standard resample-propagate procedure once more gives $p^N\left(X_{t+M}\middle|Y^{t+M}, \Theta\right)$ and completes the loop.

**Resample.** Draw a size-$N$ set of index $(k^{(i)})_{i\in\mathcal{N}}$ from $\mathcal{N}$, with weights proportional to the joint likelihood of $Y_{t+M} = \left(Y_{t+M}^H, Y_{t+M}^L\right)$

$$w_{t+M|t+M-1}^{(i)} \propto p\left(Y_{t+M}\middle|\widetilde{X}_{t+1:t+M-1}^{(i)}, \Theta\right), \tag{11}$$

where $k^{(i)}$ is the index from the $i$-th draw.

**Propagate.** For each draw $k^{(i)}$, draw $X_{t+M}^{(i)}$ from the conditional distribution

$$X_{t+M}^{(i)} \sim p\left(X_{t+M}\middle|\widetilde{X}_{t+1:t+M-1}^{(k^{(i)})}, Y_{t+M}, \Theta\right). \tag{12}$$

The set $\left\{X_{t+M}^{(i)}\right\}_{i\in\mathcal{N}}$ represents $N$ samples from $p\left(X_{t+M}\middle|Y_{t+M}, \Theta\right)$ and is denoted by $p^N\left(X_{t+M}\middle|Y_{t+M}, \Theta\right)$.

## 3.5 Extentions: Parameter Learning

Various particle filters have been developed by the literature to achieve more general tasks such as sequential parameter learning. To illustrate the flexibility of the smoothing-based approach to filtering, we examine the particle learning method of Carvalho et al. (2010) for MFSSMs. Similar extensions can be made to the sequential parameter learning algorithms of

Liu and West (2001), Storvik (2002), and many others.

In particle learning, the parameter prior is usually assumed to be conjugate. Under this assumption, the prior and posterior are distributions of the same type, and can be fully determined by their sufficient statistics. Sequential parameter estimation is achieved by augmenting the state space by sufficient statistics. To illustrate, we denote the set of sufficient statistics at time $t$ by $s_t$. The state-space model is fully characterized by the augmented state variable $(X, s)$. Prior to time $t+M$, particle learning takes the standard resample-propagate procedure

$$\begin{aligned} p\left(X_{t+l+1}, s_{t+l+1}, \Theta \middle| Y^{t+l+1}\right) &\propto \int p\left(Y_{t+l+1} \middle| X_{t+l}, \Theta\right) p\left(X_{t+l+1} \middle| X_{t+l}, \Theta, Y_{t+l+1}\right) \\ &\times p\left(s_{t+l+1} \middle| X_{t+l}, X_{t+l+1}, s_{t+l}, Y_{t+l+1}\right) \\ &\times p\left(\Theta \middle| s_{t+l+1}\right) dp\left(X_{t+l}, s_{t+l}, \Theta \middle| Y^{t+l}\right), \end{aligned} \tag{13}$$

where $l = 0, ..., M-2$. $p\left(Y_{t+l+1} \middle| X_{t+l}, \Theta\right)$ is the predictive likelihood, $p\left(X_{t+l+1} \middle| X_{t+l}, \Theta, Y_{t+l+1}\right)$ is the conditional state density, and $p\left(s_{t+l+1} \middle| X_{t+l}, X_{t+l+1}, s_{t+l}, Y_{t+l+1}\right)$ describes the sufficient statistics updating equation. The dynamic updating of parameter posteriors, $p\left(\Theta \middle| s_{t+l+1}\right)$, is tracked by sufficient statistics. Similar to particle filters, the smoothing state path $X_{t+1:t+M-1}$ is needed to evaluate the predictive likelihood and the conditional state density. We follow Carvalho et al. (2010) and employ the backward smoother for particle learning

$$\begin{aligned} p\left(X_{t+1:t+M-1}, s_{t-1}, \Theta \middle| Y^{t+M-1}\right) &= p\left(X_{t+M-1}, s_{t+M-1}, \Theta \middle| Y^{t+M-1}\right) \\ &\times \prod_{l=1}^{M-2} p\left(X_{t+l} \middle| X_{t+l+1}, Y^{t+l}, \Theta\right), \end{aligned} \tag{14}$$

with

$$p\left(X_{t+l} \middle| X_{t+l+1}, Y^{t+l}, \Theta\right) \propto p\left(X_{t+l+1} \middle| X_{t+l}, \Theta\right) p\left(X_{t+l} \middle| Y^{t+l}, \Theta\right), \quad l = M-2, ..., 1. \tag{15}$$

We assume now that the empirical filtering distribution $p^N\left(X_{t+l}, s_{t+l}, \Theta \middle| Y^{t+l}\right)$, $l = 1, ..., M-$

1, has been obtained from the forward filter. By definition, each of them is characterized by a sample set $(X_{t+l}, s_{t+l}, \Theta)^{(i)}_{i\in\mathcal{N}}$. Carvalho et al. (2010) mistakenly argue that the backward smoother works trivially in the same way as for particle filters. More precisely, they argue that the $i$-th sample of the smoothing distribution of $X_{t+l}$ shall be drawn from the marginal distribution $p^N\left(X_{t+l}\middle|Y^{t+l}\right)$. However, according to equation (15), smoothing shall be based on $p\left(X_{t+l}\middle|Y^{t+l}, \Theta^{(i)}\right)$, which can be obtained only through refiltering for each fixed parameter draw, $\Theta^{(i)}$. To address the issue of smoothing, Yang et al. (2018) link the unknown but necessary distribution, $p\left(X_{t+l}\middle|Y^{t+l}, \Theta\right)$, to $p\left(X_{t+l}, \Theta\middle|Y^{t+l}\right)$, which is available, through a multivariate normal approximation. Our view is that by directly employing MFPF for $p\left(X_{t+l}\middle|Y^{t+l}, \Theta^{(i)}\right)$, we avoid the error stemmed from multivariate normal approximations.

# 4 Empirical Applications: Asset Return Predictability

As documented in previous literature, macroeconomic survey data can predict key financial variables such as asset returns and yield changes; see, for instance, Chun (2010), Chernov and Mueller (2012), Eriksen (2017), and Foroni et al. (2018). In this section, we apply MFPF to predictive regressions in which macroeconomic survey data are used to forecast the returns of Treasury bonds and US trade-weighted dollar index. Survey data are forward-looking and have been shown to have additional explanatory power that is empirically not reflected in traditional predictors. Survey data are only quarterly available, and the usual practice is to temporally aggregate all variables to a quarterly frequency. However, asset returns exhibit time-varying volatility, which can become weak at lower frequencies such as quarterly frequencies. Ignoring time-varying volatility might give biased model inference. In contrast, a monthly-evolving model can preserve the high-frequency nature of the volatility dynamics and allows for volatility timing, which is economically important as empirically shown by Moreira and Muir (2017). In practice, forecasting asset returns at a monthly frequency is also common for practitioners who manage portfolios and volatility at a relatively high frequency such as monthly. Thus, we are aiming at incorporating quarterly survey variables into monthly-evolving predictive regressions. This section will conduct a comprehensive specification analysis and evaluate the performance of various model specifications. In particular, we discuss the predictive power of quarterly predictors and the role of stochastic volatility.

## 4.1 Model Specification

We use the Survey of Professional Forecasters (SPF) of the growth rates of industrial production (IP) and real personal consumption expenditures (PCE), and CPI inflation rates (CPI) for the next quarter as the predictors. The complete dataset is provided by the Philadelphia Fed and dates back to 1968:Q4. However, it has been handled by the Philadelphia Fed and released in real time only since 1990:Q2. Thus, we only consider the sample period 1990:Q2 to 2018:Q3. The goal is to predict monthly returns. The quarterly survey data are reported in annualized terms. For our studies, we convert them to quarterly terms to match their sampling frequency. To incorporate these quarterly predictors, we assume that each of them, denoted by $Z$, is the quarterly sum of a monthly-evolving linear-Gaussian process $X$

$$\begin{aligned} Z_{t+3} &= X_{t+1} + X_{t+2} + X_{t+3}, \\ X_{t+1} &= K_{X,0} + K_{X,1} X_t + \sigma_X \epsilon^X_{t+1}, \end{aligned} \tag{16}$$

where $\epsilon^X_{t+1}$ is i.i.d. standard normal. Economically, $X$ can be understood as the monthly forecast implied by the quarterly survey. Using a linear-Gaussian model is consistent with the stylized empirical fact that the survey variables are mean-reverting and persistent. To forecast asset returns in the next month, we use $X$ instead of the quarterly observations directly, i.e.,

$$r_{t+1} = K_{r,0} + K_{r,1} X_t + \sigma_r \epsilon^r_{t+1}, \tag{17}$$

where $r_{t+1}$ is the monthly return and $K_r = (K_{r,0}, K_{r,1})$ is the prediction coefficient. For Treasury bonds, $r_{t+1}$ is the average logarithmic return of 5-, 10-, 15-, 20-, 25-, and 30-year bonds in excess of the 1-month bond return. The bond returns are obtained from the yield curve data from Gürkaynak et al. (2007). For the dollar index, $r_{t+1}$ is the monthly logarithmic return and obtained from Federal Reserve Economic Data. $Z$ is observed only at the end of each quarter, whereas $X$ is not observable. Casting the predictive regressions in state-space form, $Z$ and $r$ play the roles of low- and high-frequency observations, respectively. Predicting the time-$(t+1)$ return exploits the time-$t$ filtering density of $X$, which can be obtained from MFPF. The forecasting error is assumed to be i.i.d. normal, with a standard deviation of $\sigma_r$. We assume all parameters in equations (16) and (17) to be constant. We call this model the mixed-frequency constant-variance (MF-CV) model. Incorporating mixed-frequency data by a simple aggregation scheme as in equation (16) is common in the macroeconomic forecasting literature and has been adopted, among others, by Mariano and Murasawa (2003), Aruoba et al. (2009), Schorfheide and Song (2015), Marcellino et al. (2016), Schorfheide et al. (2018), and Leippold and Yang (2018).

There is rich evidence of time variation in asset return volatility at monthly or higher frequencies. Thus, we also consider mixed-frequency predictive regressions with stochastic volatility (MF-SV)

$$\begin{aligned} r_{t+1} &= K_{r,0} + K_{r,1} X_t + \sqrt{V_t}\epsilon^r_{t+1}, \\ \ln V_{t+1} &= K_{V,0} + K_{V,1} \ln V_t + \sigma_V \epsilon^V_{t+1}. \end{aligned} \tag{18}$$

The above volatility specification follows Johannes et al. (2014). The instantaneous variance is assumed to be a log-linear Gaussian process, with the volatility leverage effect captured by a correlation between shocks to index returns and variance $\rho = \mathrm{corr}(\epsilon^r_{t+1}, \epsilon^V_{t+1})$. The log-linear specification has several advantages compared to other volatility models considered in the literature. First of all, it is the most parsimonious specification to model mean reversion and persistence of the volatility dynamics. Second, it is able to generate sufficient skewness and kurtosis as empirically observed in the data when compared to other volatility models such as the model of Heston (1993). Third, it guarantees that the instantaneous variance remains strictly positive, which is not satisfied by many other models in a discrete-time setup.

To evaluate the relative performance of mixed-frequency models, we also consider quarterly aggregate models. At the end of each quarter, these models use the survey data to forecast the next quarter's return

$$\begin{aligned} r_{t+1,t+3} &= r_{t+1} + r_{t+2} + r_{t+3}, \\ r_{t+1,t+3} &= K_{r,0} + K_{r,1} Z_t + \sigma_r \epsilon^r_{t+1,t+3}. \end{aligned} \tag{19}$$

To examine whether incorporating stochastic volatility improves the forecasting power at a quarterly frequency, we also consider quarterly aggregate models with stochastic volatility, where the instantaneous volatility is assumed to take the same form as equation (18) but evolves at a quarterly frequency. We call these models Q-CV and Q-SV models, respectively.

## 4.2 Model Estimation

The above SV models are based on a nonlinear state-space system that does not have a closed-form solution to the associated filtering problem. Hence, the Kalman filter is not applicable. We use an MCMC iterator to make inference about $\Theta$. The inference procedure relies on the Bayes' rule

$$p\left(\Theta \middle| Y^T\right) \propto p\left(Y^T \middle| \Theta\right) p(\Theta), \tag{20}$$

where $p(\Theta)$ is the parameter prior. Equation (20) states that the parameter posterior is available if one can evaluate the joint likelihood of all observations for each choice of $\Theta$. To this end, we use MFPF to compute the joint likelihood, which can be further factorized into the conditional likelihood of observations in each period

$$p\left(Y^T\middle|\Theta\right) \;=\; \prod_{t=1}^{T} p\left(Y_t\middle|Y_{t-1},\Theta\right). \tag{21}$$

The conditional likelihood can be evaluated sequentially given the filtering distribution obtained from MFPF. For example, for each $M$ units of time, we have

$$\begin{aligned} p\left(Y_{t+l+1}\middle|Y_{t+l},\Theta\right) &\approx \int p\left(Y_{t+l+1}\middle|X_{t+l},\Theta\right) dp^N\left(X_{t+l}\middle|Y^{t+l},\Theta\right), \quad l=0,...,M-2, \\ p\left(Y_{t+M}\middle|Y_{t+M-1},\Theta\right) &\approx \int p\left(Y_{t+M}\middle|X_{t+1:t+M-1},\Theta\right) dp^N\left(X_{t+1:t+M-1}\middle|Y^{t+M-1},\Theta\right). \end{aligned} \tag{22}$$

Given the approximate likelihood, we then employ a random-walk Metropolis-Hasting algorithm to draw sample parameters from the posterior. Andrieu et al. (2010) show that using the approximate likelihood instead of the exact likelihood still delivers draws from the actual posterior as the Markov chain evolves. Further, the prior is assumed to be flat, namely $p(\Theta) \propto 1$ on its support. Both assumptions simplify the sampling procedure to the greatest extent.

### 4.3 Empirical Results

Table 1 displays the estimation results when quarterly survey variables are used to predict monthly bond returns. For both CV and SV models, the estimates for the predictor dynamics are similar. However, incorporating stochastic volatility significantly affects the estimates of the prediction coefficients, $K_{r,0}$ and $K_{r,1}$, and hence the return forecasts. Moreover, the time variation and predictability in the volatility dynamics is statistically significant, as indicated by the estimate of $K_{V,1}$. We use the average monthly logarithmic likelihood, Akaike information criteria (AIC), and the $R^2$-value to measure the forecasting performance. The model likelihood can also be viewed as the density forecast of observations.

All evidence points to the fact that SV models empirically outperform CV models in predicting monthly bond returns. First of all, for all predictors, SV models deliver a larger likelihood, which is further translated, by adjusting the number of parameters, into a smaller value of AIC. Second, the prediction $R^2$ becomes larger when stochastic volatility is incorporated. For each predictor, the $R^2$ increases from 2.32%, 0.00%, and 0.99% to 2.84%, 0.31%, and 1.35%, respectively. There is no return predictability by PCE at a monthly frequency, but incorporating stochastic volatility raises the forecasting power slightly. We can draw two conclusions. First, quarterly survey variables predict monthly bond returns, which justifies the use of mixed-frequency models when forecasting monthly returns is the objective. Second, mixed-frequency models preserve the high-frequency evolution of volatility, which in return improves return predictability.

Next, we examine the performance of temporally aggregate models. Table 2 displays the estimation results for Q-CV and Q-SV models. At a quarterly frequency, there is reasonable predictability by all three survey variables, with an $R^2$ of 4.27%, 0.88%, and 1.75%. However, unlike mixed-frequency models, the improvement from stochastic volatility is weaker. First of all, the volatility becomes less persistent and thus less predictable, as indicated by the estimate of $K_{V,1}$. The reason is trivial that temporal aggregation tends to smooth out the high-frequency variation of volatility. Accordingly, there is no improvement in $R^2$, though the likelihood becomes slightly larger. AIC does not support Q-SV models, either. To summarize, at a quarterly frequency, the benefit from incorporating stochastic volatility is limited.

[Table 1 about here.]

[Table 2 about here.]

We also use survey variables to predict US dollar index returns, with the estimation results reported in Tables 3 and 4, respectively. The results are similar except for PCE. Indeed, adding stochastic volatility increases the model likelihood but results in worse $R^2$. Still, we see a stronger forecasting power for SV-IP and SV-CPI models than CV models. Turning to quarterly aggregate models, we find no evidence of improvement due to stochastic volatility. The log likelihood and $R^2$ are indistinguishable for Q-CV and Q-SV models. The AIC for Q-SV models become even worse for IP and PCE, suggesting that stochastic volatility does not improve the model at a quarterly frequency.

[Table 3 about here.]

[Table 4 about here.]

# 5 Conclusion

In this paper, we address the issue of state-space modeling with mixed-frequency data. To this end, we propose a general sequential particle filtering framework. The advantage of the mixed-frequency particle filter is that it allows non-Gaussian shocks and nonlinear dynamics, for instance, stochastic volatility models. In dealing with mixed-frequency observations, we employ a smoothing-technique to proceed with the filter. This smoothing mitigates the notorious issue of sample degeneracy that arises from mixed-frequency data.

As empirical applications for our mixed-frequency particle filter, we examine the forecasting power of quarterly economic surveys for Treasury bond returns and US trade-weighted dollar index. We find that monthly-evolving stochastic volatility models produce better forecasts than constant volatility models, whereas quarterly aggregate models do not. The empirical studies justify the use of mixed-frequency models that preserve the high-frequency nature of volatility dynamics.

# References

Andrieu, Christophe, Arnaud Doucet, and Roman Holenstein, 2010, Particle markov chain monte carlo methods, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72, 269–342.

Aruoba, S Borağan, Francis X Diebold, and Chiara Scotti, 2009, Real-time measurement of business conditions, *Journal of Business & Economic Statistics* 27, 417–427.

Bai, Jennie, Eric Ghysels, and Jonathan H. Wright, 2013, State space models and MIDAS regressions, *Econometric Reviews* 32, 779–813.

Carvalho, Carlos, Michael S. Johannes, Hedibert F. Lopes, and Nick Polson, 2010, Particle learning and smoothing, *Statistical Science* 25, 88–106.

Chernov, Mikhail, and Philippe Mueller, 2012, The term structure of inflation expectations, *Journal of Financial Economics* 106, 367–394.

Chun, Albert Lee, 2010, Expectations, bond yields, and monetary policy, *The Review of Financial Studies* 24, 208–247.

Creal, Drew, Bernd Schwaab, Siem Jan Koopman, and Andre Lucas, 2014, Observation-driven mixed-measurement dynamic factor models with an application to credit risk, *Review of Economics and Statistics* 96, 898–915.

Doucet, Arnaud, Simon Godsill, and Christophe Andrieu, 2000, On sequential monte carlo sampling methods for bayesian filtering, *Statistics and Computing* 10, 197–208.

Durbin, James, and Siem Jan Koopman, 2012, *Time series analysis by state space methods*, volume 38 (Oxford University Press).

Eriksen, Jonas N, 2017, Expected business conditions and bond risk premia, *Journal of Financial and Quantitative Analysis* 52, 1667–1703.

Foroni, Claudia, Pierre Guérin, and Massimiliano Marcellino, 2018, Using low frequency information for predicting high frequency variables, *International Journal of Forecasting* 34, 774–787.

Frühwirth-Schnatter, Sylvia, 1994, Data augmentation and dynamic linear models, *Journal of Time Series Analysis* 15, 183–202.

Ghysels, Eric, 2016, Macroeconomics and the reality of mixed frequency data, *Journal of Econometrics* 193, 294 – 314.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2004, The MIDAS touch: Mixed data sampling regression models, Cirano working papers, CIRANO.

Ghysels, Eric, Arthur Sinko, and Rossen Valkanov, 2007, Midas regressions: Further results and new directions, *Econometric Reviews* 26, 53–90.

Giannone, Domenico, Lucrezia Reichlin, and David Small, 2008, Nowcasting: The real-time informational content of macroeconomic data, *Journal of Monetary Economics* 55, 665–676.

Godsill, Simon J., Arnaud Doucet, and Mike West, 2004, Monte carlo smoothing for nonlinear time series, *Journal of the American Statistical Association* 99, 156–168.

Gordon, Neil J., David J. Salmond, and Adrian F.M. Smith, 1993, Novel approach to nonlinear/non-gaussian bayesian state estimation, in *IEE Proceedings F-Radar and Signal Processing*, volume 140, 107–113, IET.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright, 2007, The us treasury yield curve: 1961 to the present, *Journal of monetary Economics* 54, 2291–2304.

Harvey, Andrew C., 1989, *Forecasting, structural time series models and the Kalman filter* (Cambridge university press).

Heston, Steven L, 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency, *The Review of Financial Studies* 6, 327–343.

Johannes, Michael, Arthur Korteweg, and Nicholas Polson, 2014, Sequential learning, predictability, and optimal portfolio returns, *The Journal of Finance* 69, 611–644.

Kitagawa, Genshiro, 1996, Monte carlo filter and smoother for non-gaussian nonlinear state space models, *Journal of Computational and Graphical Statistics* 5, 1–25.

Leippold, Markus, and Hanlin Yang, 2018, Mixed-frequency predictive regressions, Working paper, University of Zurich.

Liu, Jane, and Mike West, 2001, Combined parameter and state estimation in simulation-based filtering, in *Sequential Monte Carlo Methods in Practice*, 197–223 (Springer).

Marcellino, Massimiliano, Mario Porqueddu, and Fabrizio Venditti, 2016, Short-term gdp forecasting with a mixed-frequency dynamic factor model with stochastic volatility, *Journal of Business & Economic Statistics* 34, 118–127.

Mariano, Roberto S, and Yasutomo Murasawa, 2003, A new coincident index of business cycles based on monthly and quarterly series, *Journal of Applied Econometrics* 18, 427–443.

Moreira, Alan, and Tyler Muir, 2017, Volatility-managed portfolios, *The Journal of Finance* 72, 1611–1644.

Pitt, Michael K., and Neil Shephard, 1999, Filtering via simulation: Auxiliary particle filters, *Journal of the American Statistical Association* 94, 590–599.

Schorfheide, Frank, and Dongho Song, 2015, Real-time forecasting with a mixed-frequency VAR, *Journal of Business & Economic Statistics* 33, 366–380.

Schorfheide, Frank, Dongho Song, and Amir Yaron, 2018, Identifying long-run risks: A bayesian mixed-frequency approach, *Econometrica* 86, 617–654.

Storvik, Geir, 2002, Particle filters for state-space models with the presence of unknown static parameters, *IEEE Transactions on Signal Processing* 50, 281–289.

Valle e Azevedo, Joao, Siem Jan Koopman, and António Rua, 2006, Tracking the business cycle of the euro area: A multivariate model-based bandpass filter, *Journal of Business & Economic Statistics* 24, 278–290.

Yang, Biao, Jonathan R Stroud, Gabriel Huerta, et al., 2018, Sequential monte carlo smoothing with parameter estimation, *Bayesian Analysis* 13, 1137–1161.

Table 1: Mixed-Frequency Predictive Regressions for Bond Returns

| | IP | | PCE | | CPI | |
|---|---|---|---|---|---|---|
| | CV | SV | CV | SV | CV | SV |
| $K_{r,0}$ | −0.57 | −0.83 | −0.58 | −0.09 | −0.71 | −0.43 |
| | (0.41) | (0.48) | (0.29) | (0.23) | (0.39) | (0.39) |
| $K_{r,1}$ | 4.14 | 5.66 | 6.64 | 3.57 | 5.46 | 4.71 |
| | (1.44) | (1.78) | (0.80) | (1.20) | (1.76) | (1.66) |
| $\sigma_r$ | 3.81 | | 3.82 | | 3.85 | |
| | (0.15) | | (0.22) | | (0.19) | |
| $K_{X,0}$ | 0.03 | 0.03 | 0.03 | 0.04 | 0.02 | 0.03 |
| | (0.03) | (0.01) | (0.01) | (0.01) | (0.01) | (0.00) |
| $K_{X,1}$ | 0.86 | 0.87 | 0.83 | 0.75 | 0.88 | 0.88 |
| | (0.03) | (0.03) | (0.03) | (0.03) | (0.03) | (0.02) |
| $\sigma_X$ | 0.07 | 0.07 | 0.04 | 0.04 | 0.03 | 0.03 |
| | (0.01) | (0.01) | (0.00) | (0.00) | (0.00) | (0.00) |
| $K_{V,0}$ | | −1.71 | | −1.65 | | −2.37 |
| | | (1.59) | | (1.23) | | (0.69) |
| $K_{V,1}$ | | 0.74 | | 0.75 | | 0.65 |
| | | (0.24) | | (0.13) | | (0.10) |
| $\sigma_V$ | | 0.38 | | 0.37 | | 0.40 |
| | | (0.12) | | (0.70) | | (0.09) |
| $\rho$ | | 0.72 | | 0.75 | | 0.64 |
| | | (0.17) | | (0.14) | | (0.16) |
| Log likelihood | 3.34 | 3.38 | 3.55 | 3.58 | 3.66 | 3.71 |
| AIC$\times 10^{-3}$ | −2.27 | −2.29 | −2.41 | −2.43 | −2.49 | −2.52 |
| $R^2(\%)$ | 2.32 | 2.84 | 0.00 | 0.31 | 0.99 | 1.35 |

Estimation results for predicting monthly bond returns using quarterly SPF. Bond returns are the monthly average returns of 5-, 10-, 15-, 20-, 25-, and 30-year bonds in excess of 1-month bond returns. The sample period extends from 1990:Q3 to 2018:Q3.

Table 2: Quarterly Predictive Regressions for Bond Returns

| | IP | | PCE | | CPI | |
|---|---|---|---|---|---|---|
| | CV | SV | CV | SV | CV | SV |
| $K_{r,0}$ | −1.29 | −1.32 | −0.65 | −2.42 | −1.87 | −2.90 |
| | (1.40) | (1.48) | (2.21) | (2.13) | (2.44) | (2.26) |
| $K_{r,1}$ | 3.79 | 3.39 | 3.45 | 5.92 | 5.49 | 6.74 |
| | (1.70) | (1.80) | (3.46) | (3.34) | (3.88) | (3.56) |
| $\sigma_r$ | 7.23 | | 7.35 | | 7.32 | |
| | (0.49) | | (0.49) | | (0.49) | |
| $K_{V,0}$ | | −8.08 | | −9.04 | | −8.91 |
| | | (1.87) | | (1.92) | | (1.91) |
| $K_{V,1}$ | | −0.49 | | −0.67 | | −0.64 |
| | | (0.35) | | (0.36) | | (0.35) |
| $\sigma_V$ | | 0.46 | | 0.44 | | 0.40 |
| | | (0.21) | | (0.20) | | (0.19) |
| $\rho$ | | 0.25 | | 0.45 | | 0.47 |
| | | (0.35) | | (0.34) | | (0.35) |
| Log likelihood | 1.21 | 1.23 | 1.20 | 1.21 | 1.20 | 1.21 |
| AIC$\times 10^{-3}$ | −0.27 | −0.27 | −0.26 | −0.26 | −0.27 | −0.26 |
| $R^2$(%) | 4.27 | 4.27 | 0.88 | 0.88 | 1.75 | 1.75 |

Estimation results for predicting monthly bond returns using quarterly SPF. Bond returns are quarterly and aggregated from the monthly returns. Log likelihood is the average quarterly log likelihood. AIC is the Akaike information criteria. The sample period extends from 1990:Q3 to 2018:Q3.

Table 3: Mixed-Frequency Predictive Regressions for Dollar Index Returns

| | IP | | PCE | | CPI | |
|---|---|---|---|---|---|---|
| | CV | SV | CV | SV | CV | SV |
| $K_{r,0}$ | 0.30 | 0.33 | 0.81 | 0.06 | −0.09 | −0.42 |
| | (0.12) | (0.17) | (0.08) | (0.09) | (0.17) | (0.19) |
| $K_{r,1}$ | −0.66 | −0.79 | −3.26 | −0.06 | 1.22 | 1.92 |
| | (0.45) | (0.58) | (0.19) | (0.26) | (0.78) | (0.80) |
| $\sigma_r$ | 1.23 | | 1.18 | | 1.22 | |
| | (0.05) | | (0.06) | | (0.06) | |
| $K_{X,0}$ | 0.02 | 0.03 | 0.03 | 0.04 | 0.02 | 0.01 |
| | (0.02) | (0.01) | (0.01) | (0.00) | (0.01) | (0.01) |
| $K_{X,1}$ | 0.90 | 0.89 | 0.85 | 0.79 | 0.88 | 0.92 |
| | (0.03) | (0.03) | (0.02) | (0.02) | (0.03) | (0.03) |
| $\sigma_X$ | 0.07 | 0.07 | 0.04 | 0.05 | 0.02 | 0.03 |
| | (0.01) | (0.01) | (0.00) | (0.00) | (0.00) | (0.00) |
| $K_{V,0}$ | | −4.09 | | −8.10 | | −6.72 |
| | | (1.80) | | (0.36) | | (2.13) |
| $K_{V,1}$ | | 0.54 | | 0.09 | | 0.25 |
| | | (0.20) | | (0.04) | | (0.24) |
| $\sigma_V$ | | 0.48 | | 0.54 | | 0.54 |
| | | (0.12) | | (0.10) | | (0.16) |
| $\rho$ | | 0.17 | | 0.41 | | 0.40 |
| | | (0.20) | | (0.19) | | (0.23) |
| Log likelihood | 4.48 | 4.51 | 4.71 | 4.72 | 4.81 | 4.84 |
| AIC$\times 10^{-3}$ | −3.05 | −3.06 | −3.21 | −3.21 | −3.28 | −3.29 |
| $R^2$(%) | 0.46 | 0.64 | 0.58 | 0.23 | 0.75 | 1.28 |

Estimation results for predicting dollar index returns using quarterly SPF. Dollar index returns are the monthly trade-weighted US dollar index returns. Log likelihood is the average monthly log likelihood. AIC is the Akaike information criteria. The sample period extends from 1990:Q3 to 2018:Q3.

Table 4: Quarterly Predictive Regressions for Dollar Index Returns

| | IP | | PCE | | CPI | |
|---|---|---|---|---|---|---|
| | CV | SV | CV | SV | CV | SV |
| $K_{r,0}$ | 0.76 | 0.50 | 1.11 | 0.50 | −1.11 | −1.32 |
| | (0.51) | (0.56) | (0.78) | (0.84) | (0.86) | (0.76) |
| $K_{r,1}$ | −0.39 | −0.14 | −1.04 | −0.33 | 2.64 | 2.80 |
| | (0.61) | (0.66) | (1.23) | (1.30) | (1.37) | (1.22) |
| $\sigma_r$ | 2.62 | | 2.61 | | 2.58 | |
| | (0.18) | | (0.18) | | (0.17) | |
| $K_{V,0}$ | | −14.45 | | −14.02 | | −14.66 |
| | | (3.24) | | (3.19) | | (2.63) |
| $K_{V,1}$ | | −0.97 | | −0.92 | | −0.97 |
| | | (0.44) | | (0.44) | | (0.36) |
| $\sigma_V$ | | 0.19 | | 0.23 | | 0.10 |
| | | (0.18) | | (0.18) | | (0.16) |
| $\rho$ | | 0.96 | | 0.98 | | 0.98 |
| | | (0.45) | | (0.46) | | (0.43) |
| Log likelihood | 2.23 | 2.23 | 2.23 | 2.23 | 2.24 | 2.24 |
| AIC$\times 10^{-3}$ | −0.50 | −0.49 | −0.50 | −0.49 | −0.50 | −0.50 |
| $R^2(\%)$ | 0.36 | 0.36 | 0.64 | 0.64 | 3.22 | 3.22 |

Estimation results for predicting quarterly trade-weighted US dollar index returns using quarterly SPF. Dollar index returns are the quarterly trade-weighted US dollar index returns. Log likelihood is the average quarterly log likelihood. AIC is the Akaike information criteria. The sample period extends from 1990:Q3 to 2018:Q3.