

Discussion Paper

Deutsche Bundesbank
No 25/2023

**Precision-based sampling for state space
models that have no measurement error**

Elmar Mertens

Editorial Board:

Daniel Foos
Stephan Jank
Thomas Kick
Martin Kliem
Malte Knüppel
Christoph Memmel
Hannah Paule-Paludkiewicz

Deutsche Bundesbank, Wilhelm-Epstein-Straße 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-0

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet <http://www.bundesbank.de>

Reproduction permitted only if source is stated.

ISBN 978-3-95729-956-7

ISSN 2941-7503

Non-technical summary

Research Question

Linear Gaussian state space models are embedded in many modern time series models. Model estimation with Bayesian methods typically requires sampling from the state space and can be computationally demanding. Recent advances in theory and computational routines for handling large matrices have led to more efficient routines that operate on the so-called precision matrix. However, Bayesian updating of the precision matrix is ill-defined when the observables are modeled as a linear combination of the states without measurement error.

Contribution

We extend precision-based sampling to include the case when observables are modeled as a linear combination of the states. Relevant applications include trend-cycle decompositions with persistent cycles and (mixed-frequency) VARs with missing variables.

Results

The proposed sampler is considerably faster than conventional routines based on Kalman filtering and smoothing. For typical model sizes, the execution time for state space sampling is reduced by two thirds and more. Illustrated with an application to extract a measure of trend inflation from about a dozen of input variables, model estimation takes just about 2 instead of 13 hours.

Nichttechnische Zusammenfassung

Fragestellung

Lineare Zustandsraummodelle mit normalverteilten Schocks sind Bestandteil vieler empirischer Zeitreihenmodelle. Die Modellschätzung erfolgt häufig mit bayesianischen Methoden, welche die Simulation des Zustandsraummodells erfordern, was rechentechnisch aufwendig sein kann. Effiziente Rechenroutinen nutzen oft die sogenannte Präzisionsmatrix, um die Schätzung anspruchsvoller Modelle meistern zu können. Wenn die beobachteten Zeitreihenwerte als lineare Kombinationen der Zustände ohne Messfehler modelliert werden, ist eine bayesianische Verarbeitung der Präzisionsmatrix allerdings nicht vollständig definiert.

Beitrag

Wir erweitern die Anwendbarkeit von Methoden, die auf der Präzisionsmatrix basieren, auf den Fall, in dem die beobachteten Zeitreihenwerte als lineare Kombinationen der Zustände ohne Messfehler modelliert werden. Zu den relevanten Anwendungen gehören Trend-Zyklus Modelle und Vektorautoregressionen mit fehlenden Datenpunkten.

Ergebnisse

Die hier entwickelte Methode reduziert deutlich die erforderlichen Rechenzeiten im Vergleich zu üblichen Verfahren, welche auf dem Kalman Filter basieren. Für gängige Modellgrößen verringert sich die Verarbeitungszeitung um zwei Drittel und mehr. Bei einer Schätzung des Inflationstrends anhand eines guten Dutzends von Indikatoren benötigt die Schätzung mit der neuen Methode 2 Stunden, während das übliche Verfahren 13 Stunden benötigen würde.

Precision-based sampling for state space models that have no measurement error*

Elmar Mertens¹

¹*Deutsche Bundesbank*

Abstract

This article presents a computationally efficient approach to sample from Gaussian state space models. The method is an instance of precision-based sampling methods that operate on the inverse variance-covariance matrix of the states (also known as precision). The novelty is to handle cases where the observables are modeled as a linear combination of the states without measurement error. In this case, the posterior variance of the states is singular and precision is ill-defined. As in other instances of precision-based sampling, computational gains are considerable. Relevant applications include trend-cycle decompositions, (mixed-frequency) VARs with missing variables and DSGE models.

Keywords: State space models, signal extraction, Kalman filter and smoother, precision-based sampling, band matrix

JEL classification: C11, C32, C51

*I would like to thank Joshua Chan, Todd Clark, seminar participants at Bundesbank, as well as Juan Francisco Rubio-Ramirez (co-editor of the Journal of Economic Dynamics and Control), and two anonymous reviewers for their particularly helpful comments. Any errors or omissions should be regarded as solely those of the author. This paper and accompanying materials do not necessarily represent the views of the Deutsche Bundesbank or the Eurosystem. Declarations of interest: none. A supplementary online appendix and replication codes are available at <https://www.bundesbank.de/dp-en/2023-25-appendix> and <https://github.com/elarmertens/ABCprecisionsampler>, respectively.

1 Introduction

Many modern time-series models feature non-linear dynamics and departures from normally distributed shocks. Nevertheless, often these models have a hierarchical structure build around a state space that is conditionally linear and Gaussian (conditional on parameters that may be stochastic and time-varying).¹ When estimation is done via Bayesian MCMC methods, the need arises to efficiently sample the states.

A well-known approach to sample from a linear Gaussian state space is to perform Kalman filtering and smoothing recursions. However, for larger models these recursions come with non-negligible computational cost. Building on work by [Chib and Jeliazkov \(2006\)](#) and advances in processing sparse and banded matrices, [Chan and Jeliazkov \(2009\)](#) proposed “precision-based sampling” algorithms that solve the same problem as Kalman-based routines but with higher computational efficiency. Because of their computational advantages, precision-based sampling methods have increased in popularity, including the applications in [Eckert, Kronenberg, Mikosch, and Neuwirth \(2020\)](#), [Antolín-Díaz, Drechsel, and Petrella \(2021\)](#), [Zaman \(2021\)](#) and [Carriero, et al. \(2022\)](#), with extensions in [Hauber and Schumacher \(2021\)](#) and [Chan, Poon, and Zhu \(2023\)](#). The computational benefits of precision-based samplers have also been reviewed by [McCausland, Miller, and Pelletier \(2011\)](#).

Precision-based sampling is typically built around the fact that the updating equation for the posterior precision is linear in the prior precision, which neatly facilitates the computation. But, for this linear updating to hold, the state space needs to be formulated such that each observable variable is a linear combination of states plus an independent measurement error. Many applications fall into this class of models (for example factor models), but other applications feature measurement vectors that are exact linear combination of the states (and have no measurement error). This paper presents a precision-based sampling algorithm applicable to those cases, which include:

- Trend-cycle decompositions with persistent cycles, as in [Mertens \(2016b\)](#), and [Del Negro, Giannone, Giannoni, and Tambalotti \(2017; 2019\)](#).²
- VARs with missing values, for example due to mixed-frequencies, as in [Schorfheide and Song \(2015\)](#), and [Chan, Poon, and Zhu \(2023\)](#).
- Linearized DSGE models ([Smets and Wouters, 2007](#)), possibly with stochastic volatility ([Diebold, Schorfheide, and Shin, 2017](#)) and fat-tailed shocks ([Cúrdia, Del Negro, and Greenwald, 2014](#)).³

The precision matrix is the inverse of the variance-covariance matrix of the model’s state variables. However, absent measurement errors, a linear combination of the states is exactly pinned down by the problem’s measurement vector, which leads to a posterior

¹See, for example, [Primiceri \(2005\)](#), [Carriero, Clark, Marcellino, and Mertens \(2022\)](#) and references therein.

²When the “cycle” term of a trend model is white noise, it is straightforward to treat as measurement error.

³Of course, researchers may also prefer to allow for measurement error in the estimation of DSGE model as discussed by [Canova \(2014\)](#) and [Herbst and Schorfheide \(2014\)](#). Section 3.3 provides further discussion.

variance-covariance matrix of the state vector that is singular, and an ill-defined posterior precision matrix. As a result, conventional approaches to precision-based sampling do not apply, when the measurement equation has no error term. The approach proposed here decomposes the state vector into the part that is exactly described by the measurement vector and its counterpart, for which strictly positive-definite uncertainty remains. The sampler then applies otherwise standard steps to efficiently sample from the latter part of the state vector. The decomposition is based on a QR factorization of the matrix of measurement loadings that exploits its banded and sparse structure.⁴

A closely related approach is described by [Chan, Poon, and Zhu \(2023\)](#) who consider the missing data problem of mixed frequency VARs. Their application of precision-based sampling also rests on separating observed from unobserved state variables, but is specialized to the missing-value structure of the mixed-frequency case, which provides a direct partitioning of the VAR vector into observed and unobserved states. The approach presented here provides a generic way to handle measurement vectors as arbitrary linear combination of the states. In a similar vein, [Grant and Chan \(2017a,b\)](#) provide a precision-based sampler for trend extraction from a univariate series, when the cycle follows an AR(2). This paper presents a generalization to the multivariate case with VAR(p) dynamics for the cycle, and nests it into a more general sampling approach.

The remainder of this paper describes the standard case of precision-based sampling in Gaussian state space models with measurement error (Section 2), and provides a set of example models with and without measurement error (Section 3), before turning to precision-based sampling without measurement error (Section 4). These sampling methods are evaluated in Section 5 with a simulation study, and illustrated in Section 6 with an application to estimate a common inflation trend from measures of realized and expected inflation based on [Mertens \(2016b\)](#). Section 7 concludes. Additional results and codes are provided in a supplementary online appendix and a GitHub repository at <https://github.com/elarmertens/ABCprecisionsampler>.

2 Precision-based sampling in the standard case

2.1 State space setup

Consider a linear state space model with N_x dimensional state vector x_t , N_y dimensional measurement vector y_t , as well as Gaussian disturbances to states and measurements, e_t and ε_t , respectively:

$$x_t = A_t x_{t-1} + B_t \varepsilon_t, \quad \varepsilon_t \sim N(0, I), \quad x_0 \sim N(\mu_0, \Sigma_0). \quad (1)$$

$$y_t = C_t x_t + D_t e_t, \quad e_t \sim N(0, I) \quad (2)$$

For sake of generality, the state-space matrices $\{A_t\}_{t=1}^T$, $\{B_t\}_{t=1}^T$, $\{C_t\}_{t=1}^T$ and $\{D_t\}_{t=1}^T$ can be time-varying but are treated as known for the purpose of sampling the states.⁵ Hence,

⁴A related, but different, use of reduced-rank methods in precision-based sampling is considered by [Chan, Eisenstat, and Strachan \(2020\)](#) in their study of large time-varying parameter models where parameter drift has a factor structure.

⁵Time variation in measurement loadings can also capture a time-varying length of the measurement vector, y_t , which might arise, for example, in missing-data problems as illustrated by some of our

for ease of exposition, we largely focus on the case of time-invariant parameters, $A_t = A$, $B_t = B$, and $D_t = D$, except for the case of time-varying measurement loadings, which naturally arises in applications with missing data. Moreover, we denote a factorization (such as the Choleski decomposition) of the prior variance of x_0 as $\Sigma_0 = B_0 B_0'$.

We consider the problem of drawing from the posterior for $\{x_t\}_{t=0}^T$ given a sequence of observations $\{y_t\}_{t=1}^T$ conditional on a prior $x_0 \sim N(\mu_0, \Sigma_0)$ with $|\Sigma_0| \neq 0$, and refer to this problem as “sampling from the state space.” A classic approach to state space sampling are Kalman filtering and smoothing, which iterate recursively over the observations $\{y_t\}_{t=0}^T$ (once forward and once backward), with details described in [Anderson and Moore \(1979\)](#), [Kailath, Sayed, and Hassibi \(2000\)](#) or [Durbin and Koopman \(2012\)](#). Alternatively, following [Chan and Jeliazkov \(2009\)](#) and [Durbin and Koopman \(2012\)](#), the system can be vectorized and cast into the form of a static signal-extraction problem:

$$\mathbf{A}\mathbf{X} = \mathbf{X}_0 + \mathbf{B}\boldsymbol{\varepsilon}, \quad \boldsymbol{\varepsilon} \sim (\mathbf{0}, \mathbf{I}), \quad (3)$$

$$\mathbf{Y} = \mathbf{C}\mathbf{X} + \mathbf{D}\mathbf{e}, \quad \mathbf{e} \sim (\mathbf{0}, \mathbf{I}), \quad (4)$$

where $\mathbf{X}$, $\mathbf{X}_0$ and $\boldsymbol{\varepsilon}$ are $\bar{N}_x \equiv N_x \cdot (T+1)$ dimensional column vectors, and $\mathbf{Y}$ and $\mathbf{e}$ are $\bar{N}_y \equiv N_y \cdot T$ vectors, that are constructed as follows:⁶

$$\mathbf{X} = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_T \end{bmatrix}, \quad \mathbf{X}_0 = \begin{bmatrix} \mu_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad \boldsymbol{\varepsilon} = \begin{bmatrix} \varepsilon_0 \\ \varepsilon_1 \\ \vdots \\ \varepsilon_T \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_T \end{bmatrix}, \quad \mathbf{e} = \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_T \end{bmatrix}, \quad (5)$$

$$\mathbf{A} = \begin{bmatrix} I & 0 & \dots & \dots & 0 \\ -A_1 & I & 0 & \dots & \vdots \\ 0 & -A_2 & I & \dots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & -A_T & I \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} B_0 & 0 & \dots & \dots & 0 \\ 0 & B_1 & 0 & \dots & \vdots \\ 0 & \dots & B_1 & \dots & \vdots \\ \vdots & \dots & \dots & \ddots & \vdots \\ 0 & \dots & \dots & 0 & B_T \end{bmatrix}, \quad (6)$$

$$\text{and } \mathbf{C} = \begin{bmatrix} C_1 & 0 & \dots & 0 \\ 0 & C_2 & \ddots & \vdots \\ \vdots & \dots & \ddots & \vdots \\ 0 & \dots & 0 & C_T \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} D_1 & 0 & \dots & 0 \\ 0 & D_2 & \ddots & \vdots \\ \vdots & \dots & \ddots & \vdots \\ 0 & \dots & 0 & D_T \end{bmatrix}. \quad (7)$$

Precision-based sampling solves the problem of drawing from $\mathbf{X}|\mathbf{Y} \sim N(\boldsymbol{\mu}, \mathbf{P}^{-1})$ given the prior $\mathbf{X} \sim N(\boldsymbol{\mu}_0, \mathbf{P}_0^{-1})$ where $\boldsymbol{\mu}_0 = \mathbf{A}^{-1}\mathbf{X}_0$ and $\mathbf{P}_0 = \mathbf{A}'(\mathbf{B}\mathbf{B}')^{-1}\mathbf{A}$ denote prior mean and prior precision, respectively. By design, sampling from $\mathbf{X}|\mathbf{Y}$ is equivalent to generating draws from a Kalman smoothing sampler. But, instead of breaking the sampling problem down into forward and backward sequences of recursive steps, precision-

applications below. For brevity, we refer to the length of y_t , however, as a constant, N_y .

⁶When the initial states are known — for example, in case of a VAR conditioned on the initial lags of the data — $\mathbf{X}$, $\mathbf{X}_0$, and $\boldsymbol{\varepsilon}$ are to be shortened, to satisfy Assumption 1 and ensure a non-degenerate distribution for $\boldsymbol{\varepsilon}$, by dropping x_0 , μ_0 and ε_0 , and setting the first element of the shortened version of $\mathbf{X}_0$ equal to $A_1 x_0$.

based sampling seeks to solve the static problem directly. While the resulting matrices are large, their banded structure allows precision-based sampling to utilize sparse matrix routines that are easily available, for example, in MATLAB, and provide more efficient implementations than the recursions of Kalman filtering and smoothing (Chan and Jeliazkov, 2009; McCausland, Miller, and Pelletier, 2011).

2.2 Assumptions on rank of variance-covariance matrices

So far, a critical assumption for the applicability of precision-based sampling has been that each observation must be afflicted by a non-degenerate measurement error while $\mathbf{e}$ and $\boldsymbol{\varepsilon}$ are independent. In addition, $\boldsymbol{\Sigma}$ needs to be non-singular. Here we state these assumptions formally:

ASSUMPTION 1 (Full-rank disturbances to the state vector) *The variance covariance matrix of disturbances to the state vector, $\boldsymbol{\Sigma} = \mathbf{B}\mathbf{B}'$ is strictly positive definite and has full rank, and thus non-singular: $|\boldsymbol{\Sigma}| \neq 0$.*

ASSUMPTION 2 (Non-zero measurement error) *The measurement error, $\mathbf{e}$, is independent from disturbances to the state vector $\mathbf{X}$, $E(\mathbf{e}\mathbf{e}') = \mathbf{0}$. Moreover, the variance-covariance matrix of the measurement errors, $\boldsymbol{\Omega} = \mathbf{D}\mathbf{D}'$ is strictly positive definite, so that $|\boldsymbol{\Omega}| \neq 0$.*

At first glance, Assumption 1 may appear restrictive; however, it can typically be satisfied by some rewriting of the state space.⁷ In contrast, Assumption 2, rules out applicability of (conventional) precision-based sampling to a number of models. Henceforth, when Assumptions 1 and 2 are satisfied, we refer to a state space system as given by (3) and (4) (in recursive form) or (3) and (4) (in static form) as a “state space with measurement error.” Finally, it is convenient to assume $\mathbf{B}$ and $\mathbf{D}$ are square so that Assumptions 1 and 2 amount $|\mathbf{B}| \neq 0$ and $|\mathbf{D}| \neq 0$:

ASSUMPTION 3 (Square shock loadings) *The shock loadings $\mathbf{B}$ and $\mathbf{D}$ are square matrices.*

Assumption 3 can be made without loss of generality, since the only alternative is a model being written with $\mathbf{B}$ and $\mathbf{D}$ having more columns than rows, which can be rewritten with the Choleski factorizations of $\mathbf{B}\mathbf{B}'$ and $\mathbf{D}\mathbf{D}'$ in lieu of the original versions of $\mathbf{B}$ and $\mathbf{D}$.

2.3 Precision-based sampling for measurement-error systems

Provided Assumptions 1 and 2 hold, standard signal extraction implies a linear updating equation for the precision matrix. This linear updating equation is an application of the

⁷For example, $|\boldsymbol{\Sigma}| = 0$ arises when the state equation reflects the companion form of a VAR(p) with $p > 1$. The supplementary online appendix illustrates how to construct a static system for this case that satisfies Assumption 1.

Woodbury matrix identity, and hinges on the invertibility of $\mathbf{D}$. The posterior moments of the state vector can then be obtained as follows:

$$\mathbf{P} = \mathbf{P}_0 + \mathbf{C}'(\mathbf{D}\mathbf{D}')^{-1}\mathbf{C} = \mathbf{A}^{*'}\mathbf{A}^* + \mathbf{C}^{*'}\mathbf{C}^*, \quad (8)$$

$$\mathbf{P}\boldsymbol{\mu} = \mathbf{P}_0\boldsymbol{\mu}_0 + \mathbf{C}'(\mathbf{D}\mathbf{D}')^{-1}\mathbf{Y} = \mathbf{A}^{*'}\mathbf{X}_0^* + \mathbf{C}^{*'}\mathbf{Y}^* \quad (9)$$

with $\mathbf{A}^* \equiv \mathbf{B}^{-1}\mathbf{A}$, $\mathbf{X}_0^* \equiv \mathbf{B}^{-1}\mathbf{X}_0$, $\mathbf{C}^* \equiv \mathbf{D}^{-1}\mathbf{C}$, and $\mathbf{Y}^* \equiv \mathbf{D}^{-1}\mathbf{Y}$.⁸

ALGORITHM 1 (Precision-based sampling for the case with measurement error)

1. Compute the Choleski factorization $\mathbf{P}^{1/2}$ of $\mathbf{P}$ using sparse matrix routines that exploit the banded structure inherited by $\mathbf{P}$ from $\mathbf{P}_0$, $\mathbf{D}$, and $\mathbf{C}$.
2. Solve $\mathbf{P}^{1/2}\mathbf{m} = \mathbf{A}^{*'}\mathbf{X}_0^* + \mathbf{C}^{*'}\mathbf{Y}^*$ for $\mathbf{m}$, which is efficiently be done by backwards substitution, since $\mathbf{P}^{1/2}$ is lower triangular. The result leads to $\mathbf{m} = \left(\mathbf{P}^{1/2}\right)' \boldsymbol{\mu}$.
3. Generate a draw from an $\bar{N}_x$ -dimensional multivariate standard-normal $\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})$.
4. Solve $\left(\mathbf{P}^{1/2}\right)' \mathbf{X} = \mathbf{m} + \mathbf{z}$ for $\mathbf{X}$, to obtain the desired draw from $p(\mathbf{X}|\mathbf{Y})$.

We refer to applications of Algorithm 1 as “conventional precision-based samplers.”

2.4 State space models without measurement error

This paper is concerned with models that have no measurement error. While retaining the state equation (3), let the measurement equation be as in (4), but with $\mathbf{e} = \mathbf{0}$

$$\mathbf{Y} = \mathbf{C}\mathbf{X}, \quad (10)$$

To ensure a meaningful signal extraction problem, we assume $\bar{N}_y < \bar{N}_x$ and that $\mathbf{C}$ has full rank. The system satisfies Assumption 1 but not Assumption 2 since there is no measurement error. In the case without measurement error, $\mathbf{Y}$ represents an exact linear combination of $\mathbf{X}$, so that the posterior variance-covariance matrix of the states is singular:

$$\mathbf{C} \text{Var}(\mathbf{X}|\mathbf{Y}) = \mathbf{0} \implies |\text{Var}(\mathbf{X}|\mathbf{Y})| = 0.$$

It follows that the posterior precision, the inverse of $\text{Var}(\mathbf{X}|\mathbf{Y})$, is ill-defined and the conventional precisions-based sampling approach described above does not apply. The next section describes some example models with and without measurement error, before Section 4 turns to a description of a sampler for the no-measurement error case, which we call the “precision-based ABC sampler.”

⁸Following a recurring theme in this paper, construction of $\mathbf{A}^*$, $\mathbf{C}^*$, $\mathbf{X}_0^*$, and $\mathbf{Y}^*$ can be performed efficiently by exploiting sparsity in $\mathbf{B}$ and $\mathbf{D}$. Invertibility of $\mathbf{B}$ and $\mathbf{D}$ is assured by Assumptions 1 and 2.

3 Some example models

This section describes a few examples of state space systems with or without measurement error. First, we describe (multivariate) trend-cycle decompositions, then VARs with missing data as in the mixed-frequency case, and finally generic state space models without measurement as they arise in linearized DSGE models.

3.1 Trend-cycle decomposition

Consider the following trend-cycle decomposition for a vector of N_y variables, y_t :

$$y_t = \Lambda \bar{y}_t + \tilde{y}_t, \quad \bar{y}_t = \bar{y}_{t-1} + \bar{B} \bar{\varepsilon}_t, \quad \tilde{y}_t = \tilde{A} \tilde{y}_{t-1} + \tilde{B} \tilde{\varepsilon}_t \quad (11)$$

and $\bar{\varepsilon}_t \sim N(0, I)$, $\tilde{\varepsilon}_t \sim N(0, I)$. The vector $\tilde{y}_t$ consists of N_y cyclical components, and $\bar{y}_t$ is a vector of $N_{\bar{y}}$ trend components. The matrix Λ can be rank deficient or not, depending on whether the model embeds cointegrating relationships or not. For concreteness, we consider the cases of no cointegration with $\Lambda = I$, henceforth called the “multivariate trend model,” and the cointegrated case of a single common trend with $\Lambda = 1$ (the unit vector), henceforth called “common trend model.” We refer to the entire model class described by (11) as “trend-cycle” models. Such trend-cycle models have been used, for example, by [Stock and Watson \(2007, 2016\)](#), [Mertens \(2016b\)](#), [Johannsen and Mertens \(2021\)](#) and [Del Negro, et al. \(2017; 2019\)](#). For ease of exposition, parameters are treated as constants and persistence in $\tilde{y}_t$ is modeled as a VAR with no more than one lag. Critically, for now, we assume shocks to trend and cycle, $\bar{\varepsilon}_t$ and $\tilde{\varepsilon}_t$ to be mutually uncorrelated, and will return to this assumption further below.

A popular application of the trend-cycle decomposition are local-level models where the cycle is serially uncorrelated, as in [Stock and Watson \(2007, 2016\)](#); in our context that means $\tilde{A} = 0$ and we let $\Lambda = I$ to obtain $y_t = \bar{y}_t + \tilde{\varepsilon}_t$. Such models directly map into state spaces with measurement error, by treating the serially uncorrelated cycle as equivalent to measurement noise. However, when the cycle is persistent ($\tilde{A} \neq 0$), the trend-cycle model does not directly map into a state space model with measurement error, at least not in the recursive form of (1) and (2).

To find a suitable representation for the trend-cycle model of (11) in the static form of a state space model with measurement error as in (3) and (4), recall the critical role of Assumption 2. Assumption 2 requires shocks to the state equation, $\mathbf{e}$ in (3), to be independent from the measurement error, $\boldsymbol{\varepsilon}$ in (4), but it does not require the elements of $\mathbf{e}$ to be mutually independent, which allows us to stack the serially correlated elements of $\{\tilde{y}_t\}_{t=1}^T$ into $\mathbf{e}$, and $\{\bar{y}_t\}_{t=1}^T$ into $\mathbf{X}$, so that $\boldsymbol{\varepsilon}$ consists of $\{\bar{\varepsilon}_t\}_{t=1}^T$.⁹ Specifically, a static representation of the trend-cycle model can be written as follows:

$$\mathbf{Y} = \mathbf{C} \bar{\mathbf{Y}} + \tilde{\mathbf{Y}}, \quad \bar{\mathbf{A}} \bar{\mathbf{Y}} = \bar{\mathbf{Y}}_0 + \bar{\mathbf{B}} \bar{\boldsymbol{\varepsilon}}, \quad \tilde{\mathbf{A}} \tilde{\mathbf{Y}} = \tilde{\mathbf{B}} \tilde{\boldsymbol{\varepsilon}}, \quad (12)$$

and $\bar{\boldsymbol{\varepsilon}} \sim N(\mathbf{0}, \mathbf{I})$, $\tilde{\boldsymbol{\varepsilon}} \sim N(\mathbf{0}, \mathbf{I})$, and assuming a prior mean of zero for the cycle. $\bar{\mathbf{A}}$, $\tilde{\mathbf{A}}$, $\bar{\mathbf{B}}$, and $\tilde{\mathbf{B}}$ are defined analogously to $\mathbf{A}$ and $\mathbf{B}$ in (6), but filled with $\bar{\mathbf{A}}$, $\tilde{\mathbf{A}}$, $\bar{\mathbf{B}}$

⁹In other words, to satisfy Assumption 2, $\mathbf{D}$ does not have to be a diagonal matrix, just square and with full rank.

and $\tilde{B}$ instead of A_t and B_t , respectively. Likewise, $\mathbf{C}$ is as defined in (7), filled with $C_t = \Lambda$. In the case of a serially correlated gap, $\text{Var}(\tilde{\mathbf{Y}})$ is not anymore diagonal, but its inverse retains a banded structure, which enables the use of efficient routines for matrix computations: $\text{Var}(\tilde{\mathbf{Y}})^{-1} \equiv \tilde{\mathbf{P}}_0 = \tilde{\mathbf{A}}'(\tilde{\mathbf{B}}\tilde{\mathbf{B}}')^{-1}\tilde{\mathbf{A}}$, with $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ defined analogously to $\mathbf{A}$ and $\mathbf{B}$ in (6). The updating equations of a precision-based sampler for this case can compactly be written as follows:

$$\bar{\mathbf{P}} = \bar{\mathbf{P}}_0 + \mathbf{C}' \tilde{\mathbf{P}}_0 \mathbf{C} \quad (13)$$

$$\bar{\mathbf{P}} \bar{\boldsymbol{\mu}} = \bar{\mathbf{P}}_0 \bar{\boldsymbol{\mu}} + \mathbf{C}' \tilde{\mathbf{P}}_0 \mathbf{Y} \quad (14)$$

with $\bar{\mathbf{P}}_0$ and $\bar{\mathbf{P}}$, as well as $\bar{\boldsymbol{\mu}}_0$ and $\bar{\boldsymbol{\mu}}$ denoting prior and posterior precisions and means of $\tilde{\mathbf{Y}}$, respectively. In case of a multivariate trend model with $\mathbf{C} = \mathbf{I}$, the posterior precision of the trend simplifies even further to the sum of prior precisions of trend and cycle: $\bar{\mathbf{P}} = \bar{\mathbf{P}}_0 + \tilde{\mathbf{P}}_0$.

A precision-based sampler based on (13) and (14) for the trend-cycle model in (12) is henceforth called “trend-cycle sampler.” Grant and Chan (2017a,b) study trend-cycle decompositions for univariate data with AR(2) dynamics in the cycle.¹⁰ The trend-cycle sampler extends their application to the multivariate (and possibly cointegrated) case. This trend-cycle sampler is, however, not equipped to handle missing values or mixed-frequency data, to which we turn next.¹¹

3.2 VAR with missing observations and mixed-frequency data

Another example is a VAR(p) model with missing values. This kind of application may arise for various reasons, like the exclusion of extreme values, as in [Carriero, et al. \(2022\)](#), or in case of mixed-frequency data as in [Schorfheide and Song \(2015\)](#), and [Chan, Poon, and Zhu \(2023\)](#), or more generally when using of data series with different start (or end) dates for their availability. For ease of exposition, we omit intercepts and consider the case of a VAR(1), which is identical to our prototypical state equation in (1). When only some elements of the VAR vector, x_t , are observed at a given point in time t , the measurement equation becomes

$$y_t = C_t x_t \quad (15)$$

In the simple case where some observations of x_t are observed whereas others are not, C_t is just a matrix of zeros and ones that selects the elements of x_t for which observations at time t are available. Mixed-frequency applications can also lead to richer cases, where, for example, every few periods a moving average of current and lagged values of x_t gets observed as an element of y_t , so that the relevant entries of C_t reflect the weights of the

¹⁰In their univariate setup, [Grant and Chan \(2017a\)](#) also allow for shocks to trend and cycle to be correlated, and we discuss such an extension for the general multivariate case further below.

¹¹As further potential limitation, the trend-cycle sampler does not draw initial values for the cycle prior to $t = 1$, which may be of interest for a researcher seeking to estimate the autoregressive parameters in $\tilde{\mathbf{A}}$. In this context, [Grant and Chan \(2017a\)](#) write that initial lags of $\tilde{y}_t$ (here: $\tilde{y}_0$, or, more generally, values for $t = 0, -1, \dots, -p + 1$) are to be used, which can, of course be obtained by an additional backwards simulation step. In contrast, the sampler proposed below allows to directly draw initial values as part of its state space sampling.

moving average.¹² In either case, the number of elements of the measurement vector, y_t , becomes time-varying.¹³ Critically, the measurement equation in (15) for the missing-value VAR lacks an error term, and the standard precision-based sampler does not apply. Of course, together with the state equation (1), the measurement equation in (15) is straightforward to rewrite as $y_t = C_t A x_{t-1} + C_t \varepsilon_t$, which features a serially uncorrelated shock term. However, the shock term is not uncorrelated with all leads and lags of x_t , which violates Assumption 2 and thus invalidates application of Algorithm 1.

To apply precision-based sampling to this case, a separate approach is needed, and it is provided by [Chan, Poon, and Zhu \(2023\)](#). Here we restate main elements of their approach in our context and notation. For ease of exposition, consider the case where some elements of $\mathbf{X}$ are observed and others are not. As discussed in Section 2.1, the multivariate distribution of $\mathbf{X}$ is straightforward to derive as $\mathbf{X} \sim N(\boldsymbol{\mu}_0, \mathbf{P}_0^{-1})$. Denote the unobserved elements by $\tilde{\mathbf{X}}_2$, so that $\mathbf{X}$ consists of all elements of $\tilde{\mathbf{X}}_2$ and $\mathbf{Y}$ (albeit not typically in that order); in any case, it is straightforward to derive the conditional distribution of $\tilde{\mathbf{X}}_2|\mathbf{Y}$ based on the multivariate normal for $\mathbf{X}$. The details get a little richer when the mixed-frequency case involves observations that are linear combinations of current and lagged elements of the VAR vector, as described by [Chan, Poon, and Zhu \(2023\)](#), but the general principle remains to partition $\mathbf{X}$ into linear combinations inside and outside the span of $\mathbf{Y}$. The sampler proposed in Section 4 generalizes this principle in a setup that comprises the missing-value and mixed-frequency cases discussed here, the trend-cycle application described above (also when trend and cycle are correlated) as well as more general (but linear) mappings from state to measurement vector discussed in the remainder of this section.

3.3 “ABC” vs “ABCD” models, including DSGE models

The missing-value case of a VAR, with state equation as in (1) and error-free measurement equation as in (15) also represents a more generic “ABC” form of models, as they arise, for example, as solution to linearized DSGE models based on methods known from [Blanchard and Kahn \(1980\)](#), [King and Watson \(1998\)](#), or [Klein \(2000\)](#).¹⁴ Likewise, the trend-cycle model of equation (11) can be cast in ABC form with

$$x_t = \underbrace{\begin{bmatrix} I & 0 \\ 0 & \tilde{A} \end{bmatrix}}_{=A} x_{t-1} + \underbrace{\begin{bmatrix} \bar{B} & 0 \\ B^* & \tilde{B} \end{bmatrix}}_{=B} \underbrace{\begin{bmatrix} \bar{\varepsilon}_t \\ \tilde{\varepsilon}_t \end{bmatrix}}_{=\varepsilon_t}, \quad y_t = \underbrace{\begin{bmatrix} \Lambda & I \end{bmatrix}}_{=C} \underbrace{\begin{bmatrix} \bar{y}_t \\ \tilde{y}_t \end{bmatrix}}_{=x_t}, \quad (16)$$

¹²In this case, a recursive representation in companion form requires to track the relevant lags of x_t as part of an augmented state vector. However, for the purpose of representing this richer version of a mixed frequency example in static form, we can also think of a measurement equation that includes the relevant k lags of x_t , $y_t = \sum_{i=0}^k C_{i,t} x_{t-i}$ and fill rows of C associated with y_t with the corresponding loadings $C_{i,t} \forall i \leq k$.

¹³Alternatively, one could let $N_y = Nx$ and encode rows of C_t and corresponding elements of y_t associated with missing observations as zeros. However, such an approach would lead to a rank-deficient loading matrix C_t .

¹⁴Solutions to linearized DSGE models typically generate state space models with time-invariant state space matrices $A_t = A$, $B_t = B$, and $C_t = C$.

where in (11) we had $B^* = 0$, which is straightforward to relax in this ABC form to accomodate correlation between shocks to trends and cycles.¹⁵

Models cast in ABC form do not satisfy Assumption 2 and Algorithm 1 does not apply. Of course, such an ABC state space model can always be rewritten to show an innovation term as in the “ABCD form” known from Fernández-Villaverde, Rubio-Ramírez, Sargent, and Watson (2007): $x_t = \mathcal{A}x_{t-1} + \mathcal{B}w_t$ and $y_t = \mathcal{C}x_{t-1} + \mathcal{D}w_t$ with $w_t \sim N(0, I)$, $\mathcal{A} = A$, $\mathcal{B} = B$, $\mathcal{C} = C A$, $\mathcal{D} = C B$.¹⁶ However, since $\mathcal{B}\mathcal{D}' \neq 0$, this ABCD form cannot simply be cast into (4) and (5) by letting $\varepsilon_t = \mathcal{B}w_t$ and $e_t = \mathcal{D}w_t$ without violating the required independence between ε and e .¹⁷

In the case of the trend-cycle model (when trend and cycle are uncorrelated) it is straightforward to identify the cycle as a component orthogonal to the trend and cast the model in form of a state space with measurement error in static form. When trend and cycle are correlated, such an orthogonal component could be derived, but may not be readily obvious from mere inspection of the model, while the sampler presented in the next section handles a general B matrix in (16).¹⁸

In case of a missing-value VAR, it is straightforward to separate observed from unobserved states and derive a precision-based sampler for that case as shown by Chan, Poon, and Zhu (2023). To handle this and other cases, the generic solution proposed in this paper operates directly on the ABC form of the state space; this makes it straightforward to blend a trend-cycle model with a missing-value problem as illustrated in Section 6.

As noted before, the ABC form also corresponds to common solutions for linearized DSGE models. Of course, researchers may also prefer to allow for measurement error in the estimation of a DSGE model as discussed by Canova (2014) and Herbst and Schorfheide (2014). The addition of measurement error could, for example, be viewed as desirable in order to capture factors that are present in the data but are otherwise not modeled. In this spirit, Chari, Kehoe, and McGrattan (2007), Den Haan and Drechsel (2021), and Inoue, Kuo, and Rossi (2020) add non-structural driving processes, sometimes called “wedges,” into model equations that are akin to measurement errors. Importantly, these approaches motivate seeing those “measurement errors” or “wedges” as being serially correlated. As discussed before, it may or may not be obvious to map the case of correlated measurement errors into a conventional state space sampler (at least not in the original form presented in Chan and Jeliazkov (2009)), and the solution is cleanly embedded in the framework offered below.¹⁹

¹⁵For ease of exposition, in equation (16), B has been written in block-triangular form. The block-triangular need only represent a reduced-form factorization of the variance-covariance matrix of disturbances to the shock vector x_t , without implying a causal link between shocks to trends and cycle.

¹⁶Compared to (1) and (2), there is a subtle difference in timing of state and noise terms in the measurement equation of the ABCD setup in Fernández-Villaverde, et al. (2007), which can be ignored for the purpose of our discussion.

¹⁷The solution approach to linearized DSGE models of Sims (2002) leads to such ABCD systems, but can also be reduced to an ABC form, when the model is amenable to the methods of Klein (2000) as shown by Mertens (2016a).

¹⁸For the sampler propose below to apply, the matrix B does not even have to be block-triangular. In the block-triangular case, a multivariate extension of the univariate model with correlated shocks to trend and cycle by Grant and Chan (2017a,b), could also be obtained as follows: Let $\mathbf{Y} = \mathbf{C}\tilde{\mathbf{Y}} + \mathbf{Y}^*$ with $\mathbf{C} = \mathbf{\Lambda} \left(\mathbf{I} + \tilde{\mathbf{A}}^{-1}\mathbf{B}^*\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{A}} \right)$, where $\mathbf{\Lambda}$ is block diagonal with each diagonal block filled with Λ , and $\tilde{\mathbf{A}}\mathbf{Y}^* = \tilde{\mathbf{B}}\tilde{\varepsilon}$. Since $\text{Cov}(\tilde{\mathbf{Y}}, \mathbf{Y}^*) = \mathbf{0}$, this fits the trend-cycle setup of (12) with $\mathbf{Y}^*$ in lieu of $\tilde{\mathbf{Y}}$.

¹⁹Having said that, the specific implementations of Smets and Wouters (2007), Inoue, Kuo, and Rossi

4 Precision-based sampling w/o measurement error

This section presents a new precision-based sampler that applies to the case when the measurement vector is an exact linear combination of the states, as in (10), so that Assumption 2 does not hold. The solution proposed here is to decompose the state vector into the sum of two components: one perfectly explained by the measurement vector, and the other being the stochastically non-degenerate component of $\mathbf{X}|\mathbf{Y}$. To do so, we apply a QR decomposition to the (transpose of the) measurement loadings, $\mathbf{C}$, and obtain

$$\mathbf{C} = \mathbf{R}\mathbf{Q} = [\mathbf{R}_1 \quad \mathbf{0}] \begin{bmatrix} \mathbf{Q}_1 \\ \mathbf{Q}_2 \end{bmatrix} = \mathbf{R}_1 \mathbf{Q}_1, \quad (17)$$

where $\mathbf{R}_1$ is a $\bar{N}_y \times \bar{N}_y$, lower-triangular matrix that has full rank, and $\mathbf{Q}$ is an orthonormal matrix, $\mathbf{Q}\mathbf{Q}' = \mathbf{I}$, with conformable partitions $\mathbf{Q}_1$ and $\mathbf{Q}_2$ of dimensions $\bar{N}_y \times \bar{N}_x$ and $(\bar{N}_x - \bar{N}_y) \times \bar{N}_x$, respectively. To derive an efficient sampler for $\mathbf{X}|\mathbf{Y}$, we work with the rotated state vector $\tilde{\mathbf{X}} \equiv \mathbf{Q}\mathbf{X}$ which can be partitioned into $\tilde{\mathbf{X}}_1 \equiv \mathbf{Q}_1\mathbf{X}$ and $\tilde{\mathbf{X}}_2 \equiv \mathbf{Q}_2\mathbf{X}$. Since $\mathbf{X}$ is multivariate normal, so is $\tilde{\mathbf{X}}$:

$$\tilde{\mathbf{X}} \sim N(\tilde{\boldsymbol{\mu}}_0, \tilde{\mathbf{P}}^{-1}) \quad \text{with} \quad \tilde{\boldsymbol{\mu}}_0 = \mathbf{Q}\boldsymbol{\mu}_0, \quad \text{and} \quad \tilde{\mathbf{P}} = \mathbf{Q}\mathbf{P}_0\mathbf{Q}' = \begin{bmatrix} \tilde{\mathbf{P}}_{11} & \tilde{\mathbf{P}}_{12} \\ \tilde{\mathbf{P}}_{12}' & \tilde{\mathbf{P}}_{22} \end{bmatrix} \quad (18)$$

where $\boldsymbol{\mu}_0$ and $\mathbf{P}_0$ (prior mean and precision for $\mathbf{X}$), are the same as in Section 2, and the partitions $\tilde{\mathbf{P}}_{11}$, $\tilde{\mathbf{P}}_{12}$, $\tilde{\mathbf{P}}_{21}$, and $\tilde{\mathbf{P}}_{22}$ conform with $\tilde{\mathbf{X}}_1$ and $\tilde{\mathbf{X}}_2$. With $\mathbf{X} = \mathbf{Q}_1'\tilde{\mathbf{X}}_1 + \mathbf{Q}_2'\tilde{\mathbf{X}}_2$, we construct posterior draws of $\mathbf{X}|\mathbf{Y}$ from $\tilde{\mathbf{X}}_1|\mathbf{Y}$ and $\tilde{\mathbf{X}}_2|\mathbf{Y}$. The posterior density of $\tilde{\mathbf{X}}_1|\mathbf{Y}$ is a point mass:

$$\mathbf{Y} = \mathbf{C}\mathbf{X} = \mathbf{R}_1\tilde{\mathbf{X}}_1 \quad \Leftrightarrow \quad \tilde{\mathbf{X}}_1 = \mathbf{R}_1^{-1}\mathbf{Y}.$$

The posterior $\tilde{\mathbf{X}}_2|\mathbf{Y}$ can be deduced from the joint distribution of $\tilde{\mathbf{X}}_1$ and $\tilde{\mathbf{X}}_2$ in (18) by constructing the conditional $\tilde{\mathbf{X}}_2|\tilde{\mathbf{X}}_1$ and evaluating it at $\tilde{\mathbf{X}}_1 = \mathbf{R}_1^{-1}\mathbf{Y}$. Since $\tilde{\mathbf{X}}_1$ and $\tilde{\mathbf{X}}_2$ are jointly multivariate normal, so is also the conditional $\tilde{\mathbf{X}}_2|\tilde{\mathbf{X}}_1$, which has the following moments:

$$\begin{aligned} E(\tilde{\mathbf{X}}_2|\tilde{\mathbf{X}}_1) &= E(\tilde{\mathbf{X}}_2) + \text{Cov}(\tilde{\mathbf{X}}_2, \tilde{\mathbf{X}}_1) \text{Var}(\tilde{\mathbf{X}}_1)^{-1} (\tilde{\mathbf{X}}_1 - E(\tilde{\mathbf{X}}_1)) \\ &= \mathbf{Q}_2\tilde{\boldsymbol{\mu}}_0 - \tilde{\mathbf{P}}_{22}^{-1}\tilde{\mathbf{P}}_{21}(\tilde{\mathbf{X}}_1 - \mathbf{Q}_1\tilde{\boldsymbol{\mu}}_0), \end{aligned} \quad (19)$$

$$\begin{aligned} \text{and} \quad \text{Var}(\tilde{\mathbf{X}}_2|\tilde{\mathbf{X}}_1) &= \text{Var}(\tilde{\mathbf{X}}_2) - \text{Cov}(\tilde{\mathbf{X}}_2, \tilde{\mathbf{X}}_1) \text{Var}(\tilde{\mathbf{X}}_1)^{-1} \text{Cov}(\tilde{\mathbf{X}}_2, \tilde{\mathbf{X}}_1)' \\ &= \tilde{\mathbf{P}}_{22}^{-1}, \end{aligned} \quad (20)$$

(2020), and Den Haan and Drechsel (2021) merely use likelihood evaluations of the state spaces in their models (be it for maximum-likelihood evaluations or as part of a Metropolis-Hastings step), without need for sampling unobserved states (including “wedges” and the like). In contrast, the examples of Cúrdia, Del Negro, and Greenwald (2014) and Diebold, Schorfheide, and Shin (2017), cited in the introduction, embed state-space sampling as part of their empirical approaches.

Putting all together, we obtain:²⁰

$$E(\mathbf{X}|\mathbf{Y}) = \boldsymbol{\mu}_0 + \left(\mathbf{Q}'_1 - \mathbf{Q}'_2 \tilde{\mathbf{P}}_{22}^{-1} \tilde{\mathbf{P}}_{21} \right) \mathbf{R}_1^{-1} (\mathbf{Y} - \mathbf{C}\boldsymbol{\mu}_0), \text{ and } \text{Var}(\mathbf{X}|\mathbf{Y}) = \mathbf{Q}'_2 \tilde{\mathbf{P}}_{22}^{-1} \mathbf{Q}_2. \quad (21)$$

Similar to the case with measurement error, described in Section 2, the posterior for $\mathbf{X}$ can be cast around an inverse variance, specifically the precision for $\tilde{\mathbf{X}}_2|\tilde{\mathbf{X}}_1$. More importantly, the posterior precision for $\tilde{\mathbf{X}}_2|\tilde{\mathbf{X}}_1$ can directly be read off from the joint prior precision for $\tilde{\mathbf{X}}_1$ and $\tilde{\mathbf{X}}_2$, which is akin to the linear updating in (8) for the measurement-error case.²¹ Draws of $\mathbf{X}|\mathbf{Y}$ can then be obtained from

$$\mathbf{X} = \boldsymbol{\mu}_0 + \mathbf{Q}'_1 \tilde{\mathbf{x}}_1 - \mathbf{Q}'_2 \left(\tilde{\mathbf{P}}_{22}^{-1/2} \right)' \left(\tilde{\mathbf{P}}_{22}^{-1/2} \tilde{\mathbf{P}}_{21} \tilde{\mathbf{x}}_1 + \mathbf{z}_2 \right), \quad (22)$$

with $\tilde{\mathbf{x}}_1 = \mathbf{R}_1^{-1} (\mathbf{Y} - \mathbf{C}\boldsymbol{\mu}_0)$ and $\mathbf{z}_2 \sim N(\mathbf{0}, \mathbf{I}_2)$ an $(\bar{N}_x - \bar{N}_y)$ dimensional multivariate normal distribution. An efficient algorithm to construct these draws is as follows:

ALGORITHM 2 (Precision-based ABC sampling)

1. Compute a (sparse) QR decomposition of $\mathbf{C}'$ and collect $\mathbf{R}_1$, $\mathbf{Q}_1$ and $\mathbf{Q}_2$ as defined in equation (17).
2. Solve $\mathbf{R}_1 \tilde{\mathbf{x}}_1 = \mathbf{Y} - \mathbf{C}\boldsymbol{\mu}_0$ for $\tilde{\mathbf{x}}_1$ (via backwards substitution).
3. Collect the partitions $\tilde{\mathbf{P}}_{21}$ and $\tilde{\mathbf{P}}_{22}$ of $\tilde{\mathbf{P}}$ by computing $\tilde{\mathbf{P}}_2 = \mathbf{Q}_2 \mathbf{A}^{*'} \mathbf{A}^* \mathbf{Q}' = \begin{bmatrix} \tilde{\mathbf{P}}_{21} & \tilde{\mathbf{P}}_{22} \end{bmatrix}$, with $\mathbf{A}^* \equiv \Sigma^{-1/2} \mathbf{A}$.
4. Compute the (sparse) Choleski decomposition $\tilde{\mathbf{P}}_{22}^{1/2}$ of $\tilde{\mathbf{P}}_{22}$.
5. Solve $\left(\tilde{\mathbf{P}}_{22}^{1/2} \right)' \mathbf{m}_2 = -\tilde{\mathbf{P}}_{21} \tilde{\mathbf{x}}_1$ for $\mathbf{m}_2$ (via backwards substitution).
6. Draw $\mathbf{z}_2 \sim N(\mathbf{0}, \mathbf{I}_2)$ from an $(\bar{N}_x - \bar{N}_y)$ dimensional multivariate normal distribution.
7. Solve $\left(\tilde{\mathbf{P}}_{22}^{1/2} \right)' \tilde{\mathbf{x}}_2 = \mathbf{m}_2 + \mathbf{z}_2$ for $\tilde{\mathbf{x}}_2$ (again, via backwards substitution).
8. Let $\mathbf{X} = \boldsymbol{\mu}_0 + \mathbf{Q}'_1 \tilde{\mathbf{x}}_1 + \mathbf{Q}'_2 \tilde{\mathbf{x}}_2$.

In a typical application, the state space matrices depend on parameters that are estimated as part of a Markov chain Monte Carlo (MCMC) sampler that needs to draw many samples from the state space, each time evaluated at different parameter draws. For some models, the measurement matrix $\mathbf{C}$ reflects definitional equations that are independent of estimated parameter values, which allows to enhance computational efficiency even

²⁰The derivation utilizes $\mathbf{X} = \mathbf{Q}'_1 \tilde{\mathbf{X}}_1 + \mathbf{Q}'_2 \tilde{\mathbf{X}}_2$ so that $E(\mathbf{X}|\mathbf{Y}) = \mathbf{Q}'_1 \tilde{\mathbf{X}}_1 + \mathbf{Q}'_2 E(\tilde{\mathbf{X}}_2|\mathbf{Y})$ and $\text{Var}(\mathbf{X}|\mathbf{Y}) = \mathbf{Q}'_2 \text{Var}(\tilde{\mathbf{X}}_2|\mathbf{Y}) \mathbf{Q}_2$ since $\tilde{\mathbf{X}}_1 = \mathbf{R}_1^{-1} \mathbf{Y}$ leads to $E(\tilde{\mathbf{X}}_1|\mathbf{Y}) = \tilde{\mathbf{X}}_1$, $\text{Var}(\tilde{\mathbf{X}}_1|\mathbf{Y}) = 0$ and $\text{Cov}(\tilde{\mathbf{X}}_1, \tilde{\mathbf{X}}_2|\mathbf{Y}) = 0$.

²¹The result in (20) is also embedded in the array algorithms for state space filtering described in Kailath, Sayed, and Hassibi (2000) and reflects a special case of results in McCausland, Miller, and Pelletier (2011).

further. In those cases, the QR factorization of $\mathbf{C}$ needs to be done only once, and values for $\mathbf{R}_1$, $\mathbf{Q}_1$, and $\mathbf{Q}_2$ can be stored for repeated use in subsequent simulations.

The QR factorization of $\mathbf{C}$ offers a generic way to separate linear combinations $\mathbf{X}$ that are exactly pinned down by the measurement vector $\mathbf{Y}$ from those that are not. In Section 3, we reviewed a few examples where that mapping could easily be identified from inspection of the state space, and it may be instructive to relate those examples to the QR factorization employed here. For example, in the trend-cycle application discussed in Section 3.1, consider its ABC representation in (16) and let $\Lambda = \mathbf{I}$ so that $\mathbf{C} = [\mathbf{I} \ \mathbf{I}]$. In that case, the QR decomposition of $\mathbf{C}$ yields $\mathbf{R} = \sqrt{2} \cdot \mathbf{I}$, $\mathbf{Q}_1 = 1/\sqrt{2} \cdot [\mathbf{I} \ \mathbf{I}]$ and $\mathbf{Q}_2 = 1/\sqrt{2} \cdot [\mathbf{I} \ -\mathbf{I}]$ so that $\tilde{\mathbf{X}}_1 = \sqrt{2} \cdot \mathbf{Y}$ and $\tilde{\mathbf{X}}_2 = (\bar{\mathbf{Y}} - \tilde{\mathbf{Y}})/\sqrt{2}$. Thus, when the ABC sampler constructs the posterior of $\tilde{\mathbf{X}}_2 | \tilde{\mathbf{X}}_1$, it closely corresponds to sampling $\bar{\mathbf{Y}} | \mathbf{Y}$ as done by the precision-based sampler specialized to the trend-cycle case as in equation (13) and (14). In the missing-value example of Section 3.2, when $\mathbf{C}_t$ consists of selected rows of the identity matrix, the QR approach leads directly to $\tilde{\mathbf{X}}_1 = \mathbf{Y}$ and $\tilde{\mathbf{X}}_2$ corresponding to the unobserved data points of the VAR. The appeal of the generic ABC approach is that it easily handles also richer examples, like a combination of the trend-cycle case with missing observations, as considered in Section 6. Moreover, in the missing-observations case, the measurement loadings need not be rows of the identity matrix, but can also encode cases where, for example, quarterly observables represent a moving average of current and lagged values of a monthly variable that remains latent.

5 Simulation study

To illustrate the computational benefits of the proposed sampler, we conduct a simulation study that compares the precision-based sampler described in Section 4, and some of its variants against the disturbance smoothing sampler of Durbin and Koopman (2002), henceforth “DK.”²² The results in this section consider samplers applied to simulated data, and with known parameters. The next section provides an application in the context of an estimated model where the sampler is embedded into an MCMC procedure.

For our simulation study, we present here results for a common trend model as described in Section 3.1, but with general a VAR(p) process for the cycle as used also by Mertens (2016b), and similar to Del Negro, et al. (2017; 2019). We consider different sizes of the number of elements of y_t , N_y , the number of lags, p , and observations, T . The supplementary online appendix provides additional simulation results for sampling from a multivariate trend model and a VAR with missing values, as described in Sections 3.1 and 3.2. Throughout, all computations are implemented in MATLAB and the results presented here were generated in multi-threaded mode on a 2.1 GHz Intel Xeon machine with 32 cores in Windows. In addition, the supplementary online appendix considers some variations in computational platforms, and finds broadly similar performance gains as reported here.²³ For each application, the simulations begin by drawing a set

²²In contrast to the original state smoothing sampler of Carter and Kohn (1994), the disturbance smoothing sampler of Durbin and Koopman (2002) operates on the measurement (instead of the state) vector’s variance-covariance matrix, which is easier to invert in our case of $N_y < N_x$. Our implementation of the DK sampler builds on Jarocinski (2015).

²³The supplementary online appendix also discusses some notable distinctions in using single- vs

of VAR parameters for coefficients and variance-covariances as well as a shock vector ϵ . For simplicity, parameter values are assumed to be identical for all t . To reliably measure execution times of the different samplers, the MATLAB function `timeit` is used.²⁴

Panels A through D of Table 1 report the execution times of variants of the precision-based sampler relative to the DK sampler, and Panel E provides absolute times for the DK sampler. Panels A and B consider the ABC sampler, both inclusive and exclusive of the time spent on the QR step and other one-off computations.²⁵ The distinction is made, since typical applications require to evaluate the QR decomposition of $\mathbf{C}$ only once so that Panel A, which reports times excluding the QR calculations, should be most relevant for gauging performance in an MCMC application such as the trend-inflation example discussed in Section 6 below. Panel C reports times for a precision-based sampler specialized to the trend-cycle case that builds on equations (13) and (14) above.

As further alternative to the ABC sampler, Panel D considers the application of a conventional precision-based sampler, described in Algorithm 1, that assumes the measurement equation has a noise term, as in (2); for brevity called “noise-augmented ABC sampler.” The sampler is applied to the same data as before, but after augmenting the ABC model with a measurement error term, whose variances are set to a minimal value of $1e-10$.²⁶ Formally, the noise-augmented ABC sampler delivers only an approximate solution to the true sampling problem (by assuming measurement error when there is truly none), but could be viewed as an appealing choice for researchers in possession of an implementation of Algorithm 1 while facing an ABC problem without noise.²⁷

Table 1 documents that the precision-based sampler provides considerable benefits over the application of DK, which are boosted further when a previously prepared factorization of $\mathbf{C}$ can be used. Unsurprisingly, execution times of the DK sampler increase with N_y , p , and T . However, while the absolute execution times of the DK sampler materially rise in those directions, the relative benefits of using a precision-based sampler increase as well, so that the absolute benefits of using the ABC sampler become even more tangible (be it inclusive or exclusive of the QR factorization of $\mathbf{C}$). For example, in the common trend model of Table 1, with 20 to 25 variables, 12 lags and 800 observations, the execution time of the precision-based sampler is about one third of the DK sampler, and with use of previously prepared QR results falls to well under a tenth. Since a single draw from the DK sampler takes multiple seconds in those cases, the absolute benefits of an ABC

multithreaded computations that can lead to some variations the relative strength of specific samplers depending on computational setup. Overall, the proposed precision-based sampling routines outperform the DK sampler across these settings.

²⁴The `timeit` function measures a median execution time after conducting multiple calls designed to average out factors such as varying CPU loads or overhead due to just-in-time compilation. For further details, see https://www.mathworks.com/help/matlab/matlab_prog/measure-performance-of-your-program.html.

²⁵The results reported for the case exclusive of QR computations also exclude time spent on the preparation of index vectors used to generate sparse matrices. These index vectors store the location of non-zero matrix elements of $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$, and can be recycled for subsequent draws as well. The results inclusive of QR computations reported in Panel B of the table include repeated creation of those index vectors as well.

²⁶We would like to thank an anonymous reviewer for suggesting this exercise.

²⁷With our choice of a low value for the volatility of measurement noise of $1e - 10$, we found draws from the “noise-augmented” precision-based sampler to be similar in distribution compared to those from the ABC sampler that samples from the correct distribution in this Monte Carlo experiment.

sampler are particularly tangible as well. For MCMC applications that require, say, 20'000 MCMC draws, gaining a single second translates to a total time saving of about 5.5 hours, as Section 6 illustrates.

Panel B and Panel C report quite similar benefits of using either the precision-based sampler for the generic ABC form as in Algorithm (2) or a version specialized to the trend-cycle case with uncorrelated trend and cycle as in equations (13) and (14). Of course, the specialized sampler ekes out some slight gains over the ABC sampler, in particular when p is small; however, in all cases the execution times differ by less than a tenth of a second.²⁸ The added versatility and generic applicability of the ABC sampler comes at no particular cost in execution times, when a one-off calculation of the QR step can be used.²⁹ Moreover, standardized code for the ABC case can be levered to cover a large class of state space models — including the trend-cycle case, but also VARs with missing values, and other examples as discussed in Section 3.

Finally, Panel D considers the noise-augmented ABC sampler; its performance is underwhelming. The noise-augmented sampler always underperforms the other precision-based alternatives reported in Panels A through C, and often by a wide margin. To understand the sources of inefficiencies in this setup, recall that the ABC setup has a state vector $\mathbf{X}$ with $N_x \times T$ elements (ignoring initial values). The measurement equation, $\mathbf{Y} = \mathbf{C} \mathbf{X}$, pins down $N_y \times T$ linear combinations of $\mathbf{X}$ so that the ABC sampler (or the trend-cycle sampler reported in Panel C) has to sample only $(N_x - N_y) \times T$ elements of $\mathbf{X}$ (or linear combinations thereof), which were denoted $\tilde{\mathbf{X}}_2$ in our formal derivation of the ABC sampler in Section 4. In contrast, the noise-augmented ABC sampler attaches noise to each of the $N_y \times T$ measurement observations and has to sample from a non-degenerate distribution of $\mathbf{X}$ of dimension $N_x \times T$. In both cases, a Choleski decomposition has to be performed for a precision matrix and a vector of random variables has to be drawn, but in the ABC case the relevant dimension for both operations is of length $(N_x - N_y) \times T$, whereas it is $N_x \times T$ in the noise-augmented case. While the noise-augmented sampler may seem to offer a practical short-cut for sampling from an ABC setup, it does unnecessarily increase the scale and computational demands of the problem.

²⁸Moreover, the implementation of the ABC sampler also delivers draws of all states, including $\tilde{y}_t$, for initial values $t = 0, -1, \dots, -p + 1$, which is factored into the results shown in Panels A and B. As discussed in Section 3, the trend-cycle sampler of equations (13) and (14) does not generate those draws, and their computations are ignored in Panel D.

²⁹Having said, depending on computational constraints, use of the more specialized routines may still be attractive for particularly time-critical application.

Table 1: Execution times for sampling states from common trend model

p	N_y									
	$T = 200$					$T = 800$				
	5	10	15	20	25	5	10	15	20	25
PANEL A: ABC-PS (excl. QR) as percentage of DK										
4	14	10	35	20	27	12	30	31	25	24
8	7	14	14	13	13	14	15	12	12	13
12	10	10	11	10	11	13	8	9	9	8
24	6	6	7	3	3	5	5	5	2	2
PANEL B: ABC-PS (incl. QR) as percentage of DK										
4	29	26	70	44	62	30	70	70	69	61
8	25	31	31	30	34	37	33	28	28	28
12	21	24	24	21	23	30	19	22	20	20
24	11	12	13	6	6	10	11	11	5	6
PANEL C: Trend-Cycle PS as percentage of DK										
4	12	9	27	16	21	13	26	27	24	21
8	9	13	12	13	12	15	13	12	12	11
12	7	9	9	9	8	10	7	9	10	8
24	4	5	5	2	2	4	5	5	2	2
PANEL D: standard PS w/noise as percentage of DK										
4	26	43	96	54	66	47	72	75	65	55
8	18	43	40	38	42	42	38	32	34	40
12	31	29	29	31	29	36	21	25	27	27
24	16	17	17	8	10	14	15	16	8	11
PANEL E: DK in seconds										
4	0.02	0.05	0.06	0.16	0.20	0.05	0.13	0.23	0.44	0.78
8	0.04	0.11	0.27	0.42	0.61	0.12	0.42	1.02	1.57	2.35
12	0.06	0.27	0.49	0.76	1.18	0.22	1.00	1.77	2.95	4.50
24	0.27	0.78	1.57	5.84	7.68	1.00	2.82	5.99	22.36	30.71

Notes: Based on simulated data. Panels A through D report execution times of precision-based samplers (PS) as percentage of the execution time of the Durbin-Koopmann's disturbance smoothing sampler (DK). Execution times (in seconds) of the DK sampler are reported in Panel E. Panels A and B report execution times for the generic ABC precision-based sampler, with Panel B reflecting the use of prepared one-off computations. Panel C provides results for a PS specialized to the trend-cycle case (and prepared one-off computations). Panel D reflects calls to a standard precision-based sampler, when called with a minimal noise term added to the measurement equation. All times were measured in MATLAB (R2021b) with the `timeit` function on an Intel(R) Xeon(R) Gold 6320 CPU @ 2.1 GHz (Windows) in multi-threaded mode (for matrix operations) using 32 threads.

6 Application: common trend inflation

The previous section characterized the performance benefits of precision-based samplers when considered in isolation, and assuming parameter values are known. However, these samplers are typically not applied on a standalone basis, but embedded in broader MCMC estimation schemes that also provide inference on model parameters feeding into some of the space matrices. This section illustrates the use of two state space samplers when embedded into an MCMC estimation of a common trend model. Specifically, we consider the trend inflation model of [Mertens \(2016b\)](#) and compare the use of a conventional DK sampler against the ABC version of a precision-based sampler developed in this paper.

The trend inflation model of [Mertens \(2016b\)](#) extracts a common trend from mixed-frequency data comprising measures of realized inflation and survey forecasts. In addition to survey data, the model combines headline measures of inflation, for both PCE and CPI, with various other inflation measures intended to strip out short-run fluctuations, such as core and trimmed metrics.³⁰ In total, the data set comprises thirteen variables, and all data has been obtained from the St. Louis Fed’s FRED database. As far as available, the data set uses monthly readings from January 1960 through April 2023. However, data availability varies markedly across the different series. Sticky inflation measures are available back to 1967, whereas readings for trimmed inflation do not begin before 1978. SPF forecasts are collected at a quarterly frequency (typically in the 2nd month of each quarter), and also with varying availability for different inflation measures and forecast horizons, with forecasts for the GDP deflator dating back the farthest (until 1968). SPF forecasts come at a quarterly frequency, whereas most measures for realized inflation are produced at a monthly frequency (except for the GDP deflator). There are missing observations in the data due limited availability of a given series over the sample, and due to some series having observations only once per quarter. A detailed list of all variables and their availability is provided in [Table 2](#).

The model is a common trend model with missing observations, combining elements of each example discussed in [Section 3](#).³¹ The data is modeled as a linear combination of a single common trend and variable-specific cyclical components. The trend is assumed to follow a random walk, whereas the cyclical components are jointly modeled as a VAR(p) process, with p set 12. The data vector, y_t , consists of all variables listed in [Table 2](#), for which a data point is available at time t . Due to the missing observations, the trend-cycle representation of [Section 3.1](#) does not directly apply. There, we assumed data on y_t was continuously available, which allowed us to separate the cyclical component from the “state” vector, even when the cycle is persistent. Here, we need to track the cyclical component again as part of a state vector, X_t , that is related to y_t via a measurement equation with time-varying loadings, C_t , that reflect the time-varying availability of data,

³⁰Compared to [Mertens \(2016b\)](#), the data set used here has been slightly revised, on account of public availability of data and newly constructed measures, such as the sticky CPI data from the Atlanta Fed.

³¹As discussed by [Mertens \(2016b\)](#), the common trend assumption is grounded in assuming that inflation rates in broad-based baskets as those considered here do not permanently diverge (whereas there can be permanent shifts in their price levels). Moreover, survey expectations are assumed to be weakly rational in the sense of not permanently diverging from underlying data, which allows for persistent (but not permanent) survey errors arising from a wide class of imperfect information models as documented by [Coibion and Gorodnichenko \(2015\)](#).

Table 2: Data on realized and expected measures of inflation

Variable	Frequency	Avail. since	Number of obs.
headline PCE	monthly	1960 Jan	759
core PCE	monthly	1960 Jan	759
headline PCE	monthly	1960 Jan	760
GDP deflator	quarterly	1960 Mar	253
trimmed PCE (Dallas Fed)	monthly	1978 Jan	543
trimmed CPI (Cleveland Fed)	monthly	1983 Jan	484
median CPI (Cleveland Fed)	monthly	1983 Jan	484
sticky CPI (headline, Atlanta Fed)	monthly	1967 Jan	676
sticky CPI (core, Atlanta Fed)	monthly	1967 Dec	665
10-year expected CPI (SPF)	quarterly	1991 Nov	126
SPF: CPI change over next four quarters	quarterly	1981 Aug	167
SPF: PCE change over next four quarters	quarterly	1981 Aug	167
SPF: GDP deflator over next four quarters	quarterly	1968 Nov	213

Notes: All data has been obtained from the St. Louis Fed’s FRED database per May 9 2023. (For realized inflation measures, seasonally adjusted series have been obtained.) Unless otherwise noted, headline, core, sticky, trimmed and median inflation are measured as monthly changes. GDP deflator inflation as quarterly change. Trimmed PCE inflation and sticky core inflation measures reflect 12-month changes. Each measure is expressed as a continuously compounded, annualized rate of inflation.

as in Section 3.2. Formally, the model is the following:

$$y_t = C_t x_t + y_0, \quad x_t = \begin{bmatrix} \bar{y}_t \\ \tilde{y}_t \end{bmatrix}, \quad \bar{y}_t = \bar{y}_{t-1} + \bar{\varepsilon}_t, \quad \tilde{y}_t = \sum_{k=1}^p \tilde{A}_k \tilde{y}_{t-k} + \tilde{\varepsilon}_t \quad (23)$$

$$C_t = \begin{bmatrix} 1 & I \end{bmatrix} \quad \bar{\varepsilon}_t \sim N(0, \bar{\sigma}_t^2), \quad \tilde{\varepsilon}_t \sim N(0, \tilde{\Sigma}) \quad (24)$$

where $\bar{y}_t$ is a scalar, $\tilde{y}_t$ is a $N_y \times 1$ column vector, and the VAR for $\tilde{y}_t$ is assumed stable (eigenvalues of its companion form matrix are inside the unit circle), and we have $N_y=13$ variables. The constant y_0 is a $N_y \times 1$ vector of intercepts that captures constant differences in inflation rates from different price baskets, or constant biases in survey expectations. The specification for C_t in (24) assumes all N_y variables are observable at time t ; in case of missing data for a given t , the corresponding rows are to be omitted from C_t .³² As in the baseline version of the model in Mertens (2016b), shocks to the trend are assumed to be affected by a scalar stochastic volatility (SV) process:

$$\log \bar{\sigma}_t^2 = \log \bar{\sigma}_{t-1}^2 + \eta_t, \quad \eta_t = N(0, \sigma_\eta^2). \quad (25)$$

Estimation can be performed with an MCMC-Gibbs sampler. The model consists of a set of model parameters and two sets of latent variables: the state vector of the

³²For the purpose of extracting a common trend, it is sufficient for C_t to have unit loadings on quarterly variables, like the change in GDP prices, for reasons discussed in Mertens (2016b).

linear state space, X_t , as well as the stochastic volatility of trend shocks, $\bar{\sigma}_t$. We collect all model parameters into the vector θ , and the paths for linear states and SV into $\mathbf{X}$ and $\bar{\sigma}$; all observed data points are collected into $\mathbf{Y}$. We use conjugate priors for all parameters (Minnesota-style normals for the VAR coefficients, and inverse-Wishart and inverse-gamma for $\tilde{\Sigma}$ and σ_η , etc.), with details provided the supplementary online appendix. The MCMC sampler iterates over the following blocks:

1. Draw linear states from $p(\mathbf{X}|\theta, \bar{\sigma}, \mathbf{Y})$ which can be done with the precision-based ABC sampler or the DK sampler.
2. Draw model parameters $p(\theta|\mathbf{X}, \bar{\sigma}, \mathbf{Y})$.
3. Draw the SV path $p(\bar{\sigma}|\mathbf{X}, \theta, \mathbf{Y})$.

Blocks 2 and 3 of this algorithm are standard, with further details provided in the supplementary online appendix.


Importantly, when using the precision-based ABC sampler in step 1, the QR decomposition of the measurement loadings $\mathbf{C}$, which is a block-diagonal matrix consisting of $\{C_t\}_{t=1}^T$ as in (7), needs to be performed only once, since C_t is independent of parameter and SV estimates. The relevant performance metrics from Table 1 for our application are thus those from Panel B of the table. Nevertheless, when measuring the execution times of the MCMC sampler reported below, we include the time spent on the one-off call to the QR decomposition (as well as other one-off preparations for constructing the sparse matrices needed for the sampler).

Panel (a) of Figure 1 compares the total execution time for the MCMC sampler under different scenarios: First, when using the precision-based ABC sampler in step 1. Second, when using the DK sampler. For both cases, computations were performed in multi-threaded mode on an Intel(R) Xeon(R) Gold 6320 CPU @ 2.1 GHz (Windows) with 32 cores.³³ In each case, 20'000 MCMC draws (thereof 10'000 burnin) were generated, and the time spent on executing the entire MCMC scheme was recorded. The results are staggering: When using the precision-based sampler, it takes just about 2 hours to generates all draws, whereas the execution time rises to 13.5 hours when using the DK sampler instead.

Subsequent panels of Figure 1 shows the estimated inflation trend and selected input series. By conditioning the trend estimates on various readings, including trimmed measures and SPF forecasts, the recent surge in realized measures of inflation is attributed largely to cyclical movements, with only a mild uptick in trend estimates for the last year. The long-run forecast for a given inflation measure is the sum of the common trend level and the variable-specific intercept $E_t y_{t+\infty} = \bar{y}_t + y_0$, and is estimated to be about 2.1% for the PCE inflation rate per April 2023.

³³Times are measured by calling the functions `tic` and `toc` before and after the MCMC sequence, respectively.

Figure 1: Inflation trend estimates



Note: Panel (a) reports execution times for MCMC estimation of the common trend model, conditioned on all variables listed in Table 2, with 20'000 MCMC draws (thereof 10'000 burnin draws). “DK” refers to use of a Durbin-Koopman sampler whereas “PS” refers to a precision-based ABC sampler. All times were measured in MATLAB (R2021b) on an Intel(R) Xeon(R) Gold 6320 CPU @ 2.1 GHz (Windows) with 32 cores. Panels (b) – (d) show selected input series and the trend estimate (posterior mean, level adjusted for each series) with 68% uncertainty bands. Grey shaded areas depict the 68% uncertainty bands for missing data points of the quarterly GDP deflator and SPF forecasts in Panels (c) and (d).

7 Conclusion

Pioneered by Chib and Jeliazkov (2006) and Chan and Jeliazkov (2009), precision-based sampling has become an increasingly popular choice for drawing states from a Gaussian state space. Compared to earlier methods, such as those of Carter and Kohn (1994) or Durbin and Koopman (2002), precision-based sampling offers considerable benefits in computation.

The novelty of the proposed approach is to handle cases when the model’s measurement equation comes without a term for measurement errors so that the vector of observables is an exact linear combination of the states. In that case, the posterior variance of the states is singular and only the prior precision is well defined, but not its posterior. Similar to other applications of precision-based sampling, the computational gains are substantial. Relevant applications include trend-cycle decompositions, linearized DSGE models, or (mixed-frequency) VARs with missing variables.

References

- Anderson, Brian D. O., and John B. Moore (1979), *Optimal Filtering*, Information and System Sciences Series, Englewood Cliffs, New Jersey: Prentice-Hall Inc.
- Antolín-Díaz, Juan, Thomas Drechsel, and Ivan Petrella (2021), “Advances in nowcasting economic activity: Secular trends, large shocks and new data,” Discussion Paper 15926, CEPR.
- Blanchard, Oliver Jean, and Charles M. Kahn (1980), “The solution of linear difference models under rational expectations,” *Econometrica*, 48, 1305–1312.
- Canova, Fabio (2014), “Bridging DSGE models and the raw data,” *Journal of Monetary Economics*, 67, 1–15, <https://doi.org/10.1016/j.jmoneco.2014.06.003>.
- Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens (2022), “Addressing COVID-19 outliers in BVARs with stochastic volatility,” *Review of Economics and Statistics*, forthcoming, https://doi.org/10.1162/rest_a_01213.
- Carter, C. K., and R. Kohn (1994), “On Gibbs sampling for state space models,” *Biometrika*, 81, 541–553, <https://doi.org/10.1093/biomet/81.3.541>.
- Chan, Joshua C. C., Eric Eisenstat, and Rodney W. Strachan (2020), “Reducing the state space dimension in a large TVP-VAR,” *Journal of Econometrics*, 218, 105–118, <https://doi.org/10.1016/j.jeconom.2019.11.006>.
- Chan, Joshua C. C., and Ivan Jeliazkov (2009), “Efficient simulation and integrated likelihood estimation in state space models,” *International Journal of Mathematical Modelling and Numerical Optimization*, 1, 101–120, <https://doi.org/10.1504/IJMMNO.2009.03009>.
- Chan, Joshua C. C., Aubrey Poon, and Dan Zhu (2023), “High-dimensional conditionally gaussian state space models with missing data,” *Journal of Econometrics*, 236, p. Article 105468, <https://doi.org/10.1016/j.jeconom.2023.05.005>.

- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan (2007), “Business cycle accounting,” *Econometrica*, 75, 781–836, <https://doi.org/10.1111/j.1468-0262.2007.00768.x>.
- Chib, Siddhartha, and Ivan Jeliazkov (2006), “Inference in semiparametric dynamic models for binary longitudinal data,” *Journal of the American Statistical Association*, 101, 685–700, <https://doi.org/10.1198/016214505000000871>.
- Coibion, Olivier, and Yuriy Gorodnichenko (2015), “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 105, 2644–78, <https://doi.org/10.1257/aer.20110306>.
- Cúrdia, Vasco, Marco Del Negro, and Daniel L. Greenwald (2014), “Rare shocks, great recessions,” *Journal of Applied Econometrics*, 29, 1031–1052, <https://doi.org/10.1002/jae.2395>.
- Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti (2017), “Safety, liquidity, and the natural rate of interest,” *Brookings Papers on Economic Activity*, 48, 235–316, <https://doi.org/10.1353/eca.2017.0003>.
- Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti (2019), “Global trends in interest rates,” *Journal of International Economics*, 118, 248–262, <https://doi.org/10.1016/j.jinteco.2019.01>.
- Den Haan, Wouter J., and Thomas Drechsel (2021), “Agnostic structural disturbances (ASDs): Detecting and reducing misspecification in empirical macroeconomic models,” *Journal of Monetary Economics*, 117, 258–277, <https://doi.org/10.1016/j.jmoneco.2020.01>.
- Diebold, Francis X., Frank Schorfheide, and Minchul Shin (2017), “Real-time forecast evaluation of dsge models with stochastic volatility,” *Journal of Econometrics*, 201, 322–332, <https://doi.org/10.1016/j.jeconom.2017.08.011>.
- Durbin, J., and S. J. Koopman (2012), *Time Series Analysis by State Space Methods*, Oxford Statistical Science Series: Oxford University Press, 2nd edition.
- Durbin, J., and S.J. Koopman (2002), “A simple and efficient simulation smoother for state space time series analysis,” *Biometrika*, 89, 603–615, <https://doi.org/10.1093/biomet/89.3.603>.
- Eckert, Florian, Philipp Kronenberg, Heiner Mikosch, and Stefan Neuwirth (2020), “Tracking economic activity with alternative high-frequency data,” KOF Working papers 20-488, KOF Swiss Economic Institute, ETH Zurich, <https://doi.org/10.3929/ethz-b-000458723>.
- Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson (2007), “ABCs (and Ds) of understanding VARs,” *American Economic Review*, 97, 1021–1026, <https://doi.org/10.1257/aer.97.3.1021>.

- Grant, Angelia L., and Joshua C. C. Chan (2017a), “A Bayesian model comparison for trend-cycle decompositions of output,” *Journal of Money, Credit and Banking*, 49, 525–552, <https://doi.org/10.1111/jmcb.12388>.
- (2017b), “Reconciling output gaps: Unobserved components model and Hodrick-Prescott filter,” *Journal of Economic Dynamics and Control*, 75, 114–121, <https://doi.org/10.1016/j.jedc.2016.12.00>.
- Hauber, Philipp, and Christian Schumacher (2021), “Precision-based sampling with missing observations: A factor model application,” Discussion Papers 11/2021, Deutsche Bundesbank.
- Herbst, Edward, and Frank Schorfheide (2014), “Bayesian inference for DSGE models,” *mimeo* Board of Governors of the Federal Reserve System.
- Inoue, Atsushi, Chun-Hung Kuo, and Barbara Rossi (2020), “Identifying the sources of model misspecification,” *Journal of Monetary Economics*, 110, 1–18, <https://doi.org/10.1016/j.jmoneco.2019.01>.
- Jarocinski, Marek (2015), “A note on implementing the Durbin and Koopman simulation smoother,” *Computational Statistics & Data Analysis*, 91, 1–3, <https://doi.org/10.1016/j.csda.2015.05.001>.
- Johannsen, Benjamin K., and Elmar Mertens (2021), “A time series model of interest rates with the effective lower bound,” *Journal of Money, Credit and Banking*, 53, 1005–1046, <https://doi.org/10.1111/jmcb.12771>.
- Kailath, Thomas, Ali H. Sayed, and Babak Hassibi (2000), *Linear Estimation*, Prentice Hall Information and System Sciences Series: Pearson Publishing.
- King, Robert G., and Mark W. Watson (1998), “The solution of singular linear difference systems under rational expectations,” *International Economic Review*, 39, 1015–1026.
- Klein, Paul (2000), “Using the generalized Schur form to solve a multivariate linear rational expectations model,” *Journal of Economic Dynamics and Control*, 24, 1405–1423.
- McCausland, William J., Shirley Miller, and Denis Pelletier (2011), “Simulation smoothing for state-space models: A computational efficiency analysis,” *Computational Statistics & Data Analysis*, 55, 199–212, <https://doi.org/10.1016/j.csda.2010.07.009>.
- Mertens, Elmar (2016a), “Linear RE models: Klein vs. Sims for Linear RE Systems,” *mimeo*.
- (2016b), “Measuring the level and uncertainty of trend inflation,” *The Review of Economics and Statistics*, 98, 950–967, https://doi.org/10.1162/REST_a_00549.
- Primiceri, Giorgio E. (2005), “Time varying structural vector autoregressions and monetary policy,” *Review of Economic Studies*, 72, 821–852, <https://doi.org/10.1111/j.1467-937X.2005.00353.x>.

- Schorfheide, Frank, and Dongho Song (2015), “Real-time forecasting with a mixed-frequency VAR,” *Journal of Business & Economic Statistics*, 33, 366–380, <https://doi.org/10.1080/07350015.2014.954707>.
- Sims, Christopher A (2002), “Solving linear rational expectations models,” *Computational Economics*, 20, 1–20.
- Smets, Frank, and Rafael Wouters (2007), “Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach,” *The American Economic Review*, 97, 586–606, <https://doi.org/10.1257/aer.97.3.586>.
- Stock, James H., and Mark W. Watson (2007), “Why has U.S. inflation become harder to forecast?” *Journal of Money, Credit and Banking*, 39, 3–33, <https://doi.org/10.1111/j.1538-4616.2007.00014.x>.
- (2016), “Core inflation and trend inflation,” *The Review of Economics and Statistics*, 98, 770–784, https://doi.org/10.1162/REST_a_00608.
- Zaman, Saeed (2021), “A unified framework to estimate macroeconomic stars,” Working Papers 21-23R, Federal Reserve Bank of Cleveland, <https://doi.org/10.26509/frbc-wp-202123r>.