Comput Econ (2015) 45:239–260
DOI 10.1007/s10614-013-9418-3

Tractable Latent State Filtering for Non-Linear DSGE Models Using a Second-Order Approximation and Pruning

Robert Kollmann

Accepted: 25 December 2013 / Published online: 1 February 2014
© Springer Science+Business Media New York 2014

Abstract This paper develops a novel approach for estimating latent state variables of Dynamic Stochastic General Equilibrium (DSGE) models that are solved using a second-order accurate approximation. I apply the Kalman filter to a state-space representation of the second-order solution based on the 'pruning' scheme of Kim et al. (J Econ Dyn Control 32:3397–3414, 2008). By contrast to particle filters, no stochastic simulations are needed for the deterministic filter here; the present method is thus much faster; in terms of estimation accuracy for latent states it is competitive with the standard particle filter. Use of the pruning scheme distinguishes the filter here from the deterministic Quadratic Kalman filter presented by Ivashchenko (Comput Econ, 43:71–82, 2014). The filter here performs well even in models with big shocks and high curvature.

R. Kollmann (✉)
European Centre for Advanced Research in Economics and Statistics (ECARES),
Université Libre de Bruxelles, 50 Av. Franklin Roosevelt,
CP 114, 1050 Brussels, Belgium
e-mail: robert_kollmann@yahoo.com
URL: http://www.robertkollmann.com

R. Kollmann
Centre for Economic Policy Research (CEPR), London, United Kingdom

R. Kollmann
Université Paris-Est, Créteil, France

R. Kollmann
Centre for Applied Macroeconomic Analysis (CAMA),
Australian National University, Canberra, Australia

R. Kollmann
Globalization and Monetary Policy Institute, The Federal Reserve
Bank of Dallas, Dallas, TX, USA

Keywords Latent state filtering · DSGE model estimation · Second-order approximation · Pruning · Quadratic Kalman filter

JEL Classification C63 · C68 · E37

1 Introduction

Dynamic Stochastic General Equilibrium (DSGE) models typically feature state variables that cannot directly be measured empirically (such as preference shocks), or for which data include measurement error. A vast literature during the past two decades has taken *linearized* DSGE models to the data, using likelihood-based methods (e.g., Smets and Wouters 2007; Negro and Schorfheide 2011). Linearity (in state variables) greatly facilitates model estimation, as it allows to use the standard Kalman filter to infer latent variables and to compute sample likelihood functions based on prediction error decompositions. However, linear approximations are inadequate for models with big shocks, and they cannot capture the effect of risk on economic decisions and welfare. Non-linear approximations are thus, for example, needed for studying asset pricing models and for welfare calculations in stochastic models. Recent applied macroeconomic research has begun to take *non-linear* DSGE models to the data. This work has mainly used particle filters, i.e. filters that infer latent states using *Monte Carlo* methods.[1] Particle filters are slow computationally, which limits their use to small models.

This paper develops a novel *deterministic* filter for estimating latent state variables of DSGE models that are solved using a *second-order* accurate approximation (as derived by Jin and Judd 2000, Sims 2000, Collard and Juillard 2001, Schmitt-Grohé and Uribe 2004, Kollmann 2004 and Lombardo and Sutherland 2007). That approximation provides the most tractable *non-linear* solution technique for medium-scale models, and has thus widely been used in macroeconomics (see Kollmann 2002 and Kollmann et al. 2011 for detailed references).When simulating second-order approximated models, it is common to use the 'pruning' scheme of Kim et al. (2008), under which second-order terms are replaced by products of the linearized solution. Unless the pruning algorithm is used, second-order approximated models often generate exploding simulated time paths. Pruning is therefore crucial for applied work based on second-order approximated models. This paper hence assumes that the pruned second-order approximated model is the *true* data generating process (DGP). The method presented here exploits the fact that the pruned system is *linear* in a state vector that consists of variables solved to second- and first-order accuracy, and of products of first-order accurate variables. The pruned system thus allows convenient closed-form determination of the one-period-ahead conditional mean and variance of the state vector. I apply the linear updating rule of the standard Kalman filter to the pruned state equation.

The filter here is much faster than particle filters, as it is not based on stochastic simulations. In Monte Carlo experiments, the present filter generates more accurate estimates of latent state variables than the standard particle filter, especially with big shocks or when the model has high curvature. The filter here is also more accurate than

[1] See Fernández-Villaverde and Rubio-Ramírez (2007) and An and Schorfheide (2007) for early applications.

a conventional Kalman filter that treats the linearized model as the true DGP.[2] Due to its high speed, the filter presented here is suited for the estimation of structural model *parameters*; a quasi-maximum likelihood procedure can be used for that purpose.

This paper is complementary to Andreasen (2012) and Ivashchenko (2014) who also develop *deterministic* filters for second-order approximated DSGE models, and show that those filters can outperform particle filters. These authors too assume linear updating rules. The filter here is closest to Ivashchenko's (2014) 'Quadratic Kalman filter' (QKF) that is also based on closed-form one-step-ahead conditional moments of the state vector—the key difference is that the QKF does not use the pruning scheme.[3] The present filter (based on pruning) performs well even in models with big shocks and high curvature—for such models the QKF may generate filtered state estimates that diverge explosively from true state variables. Such stability issues never arose for the filter proposed here, in a wide range of numerical experiments. In models with small shocks and weak curvature the filter developed here and the QKF have similar performance. The present paper is also related to Andreasen et al. (2013) who likewise derive a pruned state-space representation for second-order approximated DSGE models and show how to compute moments for pruned models; these authors develop a method of moments estimator for DSGE models, but do not present a filter for latent state variables. (I learnt about Ivashchenko 2014 and Andreasen et al. 2013 after the present research had been completed.)

2 Model Format and Filter

2.1 Model Format and Second-Order Solution

Many widely-used DSGE models can be expressed as:

$$E_t G(\Omega_{t+1}, \Omega_t, \varepsilon_{t+1}) = 0, \tag{1}$$

where E_t is the mathematical expectation conditional on date t information; $G : R^{2n+m} \to R^n$ is a function, and Ω_t is an $n \times 1$ vector of endogenous and exogenous variables known at t; ε_{t+1} is an $m \times 1$ vector of serially independent innovations to exogenous variables. In what follows, ε_t is Gaussian: $\varepsilon_t \sim N(0, \xi^2 \Sigma_\varepsilon)$, where ξ is a scalar that indexes the size of shocks. The solution of model (1) is a "policy function" $\Omega_{t+1} = F(\Omega_t, \varepsilon_{t+1}, \xi)$, such that $E_t G(F(\Omega_t, \varepsilon_{t+1}, \xi), \Omega_t, \varepsilon_{t+1}) = 0 \ \forall \Omega_t$. This paper focuses on second-order accurate model solutions, namely on second-order Taylor series expansions of the policy function around a deterministic stead state, i.e.

[2] The literature has discussed 'Extended Kalman filters', i.e. Kalman filters applied to linear approximations of non-linear models; e.g., Harvey (1989).

[3] One-step-ahead moments in the QKF are derived under the assumption that estimation error of filtered states is Gaussian. The filter here does not require that assumption. Ivashchenko (2014) also applies two other deterministic filters to second-order approximated DSGE models: a Central Difference Kalman filter (Norgaard et al. 2000) and an Unscented Kalman filter (Julier and Uhlmann 2004); these filters are based on different deterministic *numerical* integration schemes for computing one-step-ahead conditional moments (no analytical closed-form expressions). Andreasen (2012) estimates a DSGE model using a Central Difference Kalman filter.

around $\xi = 0$ and a point Ω such that $\Omega = F(\Omega, 0, 0)$. Let $\omega_t \equiv \Omega_t - \Omega$. For a $q \times 1$ column vector x whose i-th element is denoted x^i, let

$$P(x) \equiv vech(xx') = ((x^1)^2, x^1x^2, \ldots, x^1x^q, (x^2)^2, x^2x^3, \ldots, x^2x^q, \ldots, (x^{q-1})^2, x^{q-1}x^q, (x^q)^2),$$

be a column vector consisting of all squares and cross-products of the elements of x.[4] The second-order accurate model solution can be written as

$$\omega_{t+1} = F_0\xi^2 + F_1\omega_t + F_2\varepsilon_{t+1} + F_{11}P(\omega_t) + F_{12} \cdot (\omega_t \otimes \varepsilon_{t+1}) + F_{22}P(\varepsilon_{t+1}), \quad (2)$$

where F_0, F_1, F_2, F_{11}, F_{12}, F_{22} are vectors/matrices that are functions of structural model parameters, but that do not depend on ξ (Sims 2000; Schmitt-Grohé and Uribe 2004). The first-order accurate (linearized) model solution is:

$$\omega^{(1)}_{t+1} = F_1\omega^{(1)}_t + F_2\varepsilon_{t+1}. \quad (3)$$

The superscript $^{(1)}$ denotes a variable solved to first-order accuracy. It is assumed that all eigenvalues of F_1 are strictly inside the unit circle, i.e. that the linearized model is stationary.

2.2 Pruning

As discussed above, I use the 'pruning' scheme of Kim et al. (2008) under which second-order terms are replaced by products of the linearized solution–i.e. $P(\omega_t)$ and $\omega_t \otimes \varepsilon_{t+1}$ are substituted by $P(\omega^{(1)}_t)$ and $\omega^{(1)}_t \otimes \varepsilon_{t+1}$, respectively. With pruning, the solution (2) is thus replaced by:

$$\omega_{t+1} = F_0\xi^2 + vF_1\omega_t + F_2\varepsilon_{t+1} + F_{11}P(\omega^{(1)}_t) + F_{12}\omega^{(1)}_t \otimes \varepsilon_{t+1} + F_{22}P(\varepsilon_{t+1}). \quad (4)$$

Note that $P(\omega_t) = P(\omega^{(1)}_t)$ and $\omega_t \otimes \varepsilon_{t+1} = \omega^{(1)}_t \otimes \varepsilon_{t+1}$ hold, up to second-order accuracy.[5] Thus, (4) is a valid second-order accurate solution. In repeated applications of (2), third and higher-order terms of state variables appear; e.g., when ω_{t+1} is quadratic in ω_t, then ω_{t+2} is quartic in ω_t; pruning removes these higher-order terms. The motivation for pruning is that (2) has extraneous steady states (not present in the original model)–some of these steady states mark transitions to unstable behavior. Large shocks can thus move the model into an unstable region. Pruning overcomes this problem. If the first-order solution is stable, then the pruned second-order solution

[4] For a square matrix M, vech (M) is the column vector obtained by vertically stacking the elements of M that are on or below the main diagonal.

[5] $\omega_t = \omega^{(1)}_t + R^{(2)}$ where $R^{(n)}$ contains terms of order n or higher in deviations from steady state. Let ω^i_t and $\omega^{(1),i}_t$ be the i-th elements of ω_t and $\omega^{(1)}_t$, respectively. Note that $\omega^i_t\omega^j_t = (\omega^{(1),i}_t + R^{(2)})(\omega^{(1),j}_t + R^{(2)}) = \omega^{(1),i}_t\omega^{(1),j}_t + \omega^{(1),i}_t R^{(2)} + \omega^{(1),j}_t R^{(2)} + R^{(4)}$; thus, $\omega^i_t\omega^j_t = \omega^{(1),i}_t\omega^{(1),j}_t + R^{(3)}$. Up to 2nd order accuracy, $\omega^i_t\omega^j_t = \omega^{(1),i}_t\omega^{(1),j}_t$ and $P(\omega_t) = P(\omega^{(1)}_t)$ holds thus. By the same logic, $\omega_t \otimes \varepsilon_{t+1} = \omega^{(1)}_t \otimes \varepsilon_{t+1}$ holds to 2nd order accuracy. See Kollmann (2004) and Lombardo and Sutherland (2007).

(4) too is stable. The subsequent discussion assumes that the **true** DGP is given by the pruned system (3),(4).

2.3 Augmented State Equation

The law of motion of $P(\omega_t^{(1)})$ can be expressed as $P(\omega_{t+1}^{(1)}) = K_{11}P(\omega_t^{(1)}) + K_{12}\omega_t^{(1)} \otimes \varepsilon_{t+1} + K_{22}P(\varepsilon_{t+1})$, where K_{11}, K_{12}, K_{22} are matrices that are functions of F_1 and F_2. Stacking this matrix equation, as well as (3) and (4) gives the following state equation:

$$\begin{bmatrix} \omega_{t+1} \\ P\left(\omega_{t+1}^{(1)}\right) \\ \omega_{t+1}^{(1)} \end{bmatrix} = \begin{bmatrix} F_0\xi^2 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} F_1 & F_{11} & 0 \\ 0 & K_{11} & 0 \\ 0 & 0 & F_1 \end{bmatrix} \begin{bmatrix} \omega_t \\ P\left(\omega_t^{(1)}\right) \\ \omega_t^{(1)} \end{bmatrix} + \begin{bmatrix} F_2 \\ 0 \\ F_2 \end{bmatrix} \varepsilon_{t+1} + \begin{bmatrix} F_{12} \\ K_{12} \\ 0 \end{bmatrix} \left(\omega_t^{(1)} \otimes \varepsilon_{t+1}\right) + \begin{bmatrix} F_{22} \\ K_{22} \\ 0 \end{bmatrix} P(\varepsilon_{t+1}). \tag{5}$$

(5) can be written as: $Z_{t+1} = g_0 + G_1 Z_t + G_2\varepsilon_{t+1} + G_{12}(\omega_t^{(1)} \otimes \varepsilon_{t+1}) + G_{22}P(\varepsilon_{t+1})$, with $Z_{t+1} \equiv (\omega_{t+1}', P(\omega_{t+1}^{(1)})', \omega_{t+1}^{(1)\prime})'$, while g_0, G_1, G_2, G_{12} and G_{22} are the first to fifth coefficient vectors/matrices on the right-hand side of (5), respectively. Thus,

$$Z_{t+1} = G_0 + G_1 Z_t + u_{t+1}, \tag{6}$$

where $G_0 \equiv g_0 + G_{22}E(P(e_{t+1}))$, while

$$u_{t+1} \equiv G_2\varepsilon_{t+1} + G_{12}(\omega_t^{(1)} \otimes \varepsilon_{t+1}) + G_{22}[P(\varepsilon_{t+1}) - E(P(\varepsilon_{t+1}))]$$

is a serially uncorrelated, mean zero, disturbance. The (conditional) variance of u_{t+1} can easily be computed (see below). Note that u_{t+1} is non-Gaussian, as u_{t+1} depends on squares and cross-products of the elements of ε_{t+1}, and on the product of ε_{t+1} and $\omega_t^{(1)}$. (The absence of serial correlation of u_{t+1} follows from the assumption that ε_{t+1} is serially independent.)

Importantly, the state equation (6) is *linear* in the augmented state vector Z_t consisting of the second- and first-order accurate variables, and of the squares and cross-products of first-order accurate variables. [6] This allows convenient closed-form determination of the one-period-ahead conditional mean and variance of the state vector (see below).

2.4 Observation Equation

At $t = 1, \ldots, T$, the analyst observes a vector y_t of n_y variables that are linear functions of the state vector ω_t plus i.i.d. measurement error that is independent of

[6] Aruoba et al. (2012) estimate a pruned *univariate* quadratic time series model, using particle filter methods. These authors discard the term that is quadratic in ε_{t+1} on the right-hand side of (4). By contrast, the paper here allows for non-zero coefficients on second-order terms in ε_{t+1}, and it develops a deterministic filter that can be applied to *multivariate* models.

the state vector, at all leads and lags: $y_t = \gamma\omega_t + \psi_t$ where γ is an n_yxn matrix and $\psi_t \sim N(0, \Sigma_\psi)$ is an n_yx1 vector of measurement errors; Σ_ψ is a diagonal matrix. The observation equation can be written as:

$$y_t = \Gamma Z_t + \psi_t, \text{ with } \Gamma \equiv (\gamma, 0). \tag{7}$$

2.5 The Filter

Let $\Upsilon^t \equiv \{y_\tau\}_{\tau=1}^t$ be the observables known at date t; $X_{t,\tau} \equiv \mathrm{E}(X_t|\Upsilon^\tau)$ and $V^X_{t,\tau} \equiv E([X_t - X_{t,\tau}][X_t - X_{t,\tau}]'|\Upsilon^\tau)$ denote the mean and variance of the column vector X_t, given Υ^τ. The unconditional mean and variance are denoted by $E(X_t)$ and $V(X_t)$.

Given $Z_{t,t}$ and $V^Z_{t,t}$, the 1st and 2nd moments of the augmented state vector Z_t conditional on Υ^t, we can compute one-period-ahead conditional moments of Z_{t+1} using (6):

$$Z_{t+1,t} = G_0 + G_1 Z_{t,t}, \tag{8}$$

$$V^Z_{t+1,t} = G_1 V^Z_{t,t} G_1' + V^u_{t+1,t}, \text{ with} \tag{9}$$

$$\begin{aligned} V^u_{t+1,t} \equiv\ & G_2 V_\varepsilon G_2' + G_{12}(\omega^{(1)}_{t,t} \otimes \Sigma_\varepsilon) G_2' + G_2(\omega^{(1)}_{t,t} \otimes \Sigma_\varepsilon)' G_{12}' \\ & + G_{12}\left\{\left(V^{\omega^{(1)}}_{t,t} + \omega^{(1)}_{t,t}\omega^{(1)\prime}_{t,t}\right) \otimes \Sigma_\varepsilon\right\} G_{12}' + G_{22} V(P(\varepsilon_{t+1})) G_{22}' \end{aligned} \tag{10}$$

(see Appendix). Here and henceforth the parameter ξ that indexes the size of the exogenous shocks is normalized at $\xi = 1$

To generate $Z_{t+1,t+1}$, $V^Z_{t+1,t+1}$, I apply the *linear* updating equation of the standard Kalman filter (e.g., Hamilton 1994, ch. 13) to the state-space representation (6), (7):[7]

$$Z_{t+1,t+1} = Z_{t+1,t} + \phi_t \cdot (y_{t+1} - y_{t+1,t}), \text{ with } y_{t+1,t} = \Gamma Z_{t+1,t}, \tag{11}$$

$$\text{and} \quad \phi_t = V^Z_{t+1,t}\Gamma'\{\Gamma V^Z_{t+1,t}\Gamma' + \Sigma_\psi\}^{-1},$$

$$V^Z_{t+1,t+1} = V^Z_{t+1,t} - V^Z_{t+1,t}\Gamma'\{\Gamma V^Z_{t+1,t}\Gamma' + \Sigma_\psi\}^{-1}\Gamma V^Z_{t+1,t}. \tag{12}$$

The filter is started with the unconditional mean and variance of Z_0 : $Z_{0,0} = E(Z_0)$, $V^Z_{0,0} \equiv V(Z_0)$; $Z_{t+1,t+1}$ and $V^Z_{t+1,t+1}$ for $t \geq 0$ are computed by iterating on (8)–(12).[8] Henceforth, I refer to this filter as the 'KalmanQ' filter. Computer code that implements KalmanQ is available from the author.

$E(Z_0)$ and $V(Z_0)$ can be computed exactly; see the Appendix. The linear updating formula (11) would be an *exact* algorithm for computing the conditional expecta-

[7] Linear updating rules are likewise assumed by Andreasen (2012) and Ivashchenko (2014) who also develop deterministic filters for second-order approximated DSGE models (see above).

[8] It is assumed that the inverse of $\Gamma V^Z_{t+1,t}\Gamma + \Sigma_\psi$ (covariance matrix of prediction errors of observables) exists. A *sufficient* condition for this is that Σ_ψ is positive definite, as assumed in the numerical experiments below.

tion $Z_{t+1,t+1}$, if Z_{t+1} and the observables were (jointly) Gaussian, as then $Z_{t+1,t+1}$ would be a *linear* function of the data. This condition is not met in the second-order approximated model, as the disturbance u_{t+1} of the state equation (6) is non-Gaussian. However, as shown below, the KalmanQ filter closely tracks the true latent states.[9] Without Gaussianity, $Z_{t+1,t+1}$ is a non-linear function of data Υ^{t+1} : $Z_{t+1,t+1} = \phi(y_{t+1}, \Upsilon^t)$. (11) can be viewed as a linear approximation of this function: $Z_{t+1,t+1} \cong \phi_t \cdot (y_{t+1} - y)$, where y is the steady state of y_{t+1} and $\phi_t \equiv \partial\phi(y_{t+1}, \Upsilon^t)/\partial y_{t+1}|_{y_{t+1}=y}$. By the Law of Iterated Expectations, $Z_{t+1,t} = E(Z_{t+1,t+1}|\Upsilon^t)$, and thus: $Z_{t+1,t+1} - Z_{t+1,t} \cong \phi_t \cdot (y_{t+1} - y_{t+1,t})$.

When the *linearized* model is the true DGP (i.e. when $F_0 = 0$, $F_{11} = 0$, $F_{12} = 0$, $F_{22} = 0$), then the filter here is identical to the conventional linear Kalman filter, and the updating formula (11) holds *exactly*. In the presence of second-order model terms, KalmanQ is more accurate than a conventional Kalman filter that assumes that the linearized model (3) is the true DGP; see below.

2.6 Quasi-Maximum Likelihood Estimation of Model Parameters

If model *parameters* are unknown, then a quasi-maximum likelihood (QML) estimate of those parameters can be obtained by maximizing the function $L(\Upsilon^T|\theta) \equiv \sum_{t=1}^{T} \ln h(y_t|y_{t,t-1}(\theta); V^y_{t,t-1}(\theta))$, with respect to the vector of unknown parameters, θ. Here $h(y|\mu; V)$ is the multivariate normal density with mean μ and variance V. For a given θ, $y_{t,t-1}(\theta) \equiv \Gamma Z_{t,t-1}(\theta)$ is the prediction of y_t generated by KalmanQ, based on date $t-1$ information, Υ^{t-1}; $V^y_{t,t-1}(\theta) \equiv \Gamma V^Z_{t,t-1}(\theta)\Gamma' + \Sigma_\psi$ is the conditional variance of y_t, given Υ^{t-1}. Under conditions discussed in Hamilton (1994, ch. 13), the QML estimator θ^{QML}_T is asymptotically normal: $\sqrt{T}(\theta^{QML}_T - \theta_0) \rightarrow N(0, (J_2(J_1)^{-1}J_2)^{-1})$, where θ_0 is the true parameter vector and $J_1 \equiv \operatorname{plim} T^{-1}\sum_{t=1}^{T} \eta_t(\theta_0)\eta_t(\theta_0)'$, with $\eta_t(\theta_0) \equiv \partial \log(h_t(\theta_0))/\partial\theta$, $h_t(\theta) \equiv h(y_t|y_{t,t-1}(\theta_0); V^z_{t,t-1}(\theta_0))$ and $J_2 \equiv \operatorname{plim} T^{-1}\sum_{t=1}^{T} \partial^2 \log(h_t(\theta_0))/\partial\theta\partial\theta'$.

2.7 Ivashchenko's (2014) Quadratic Kalman Filter (QKF)

The QKF posits that the *unpruned* second-order approximated model (2) is the true DGP [instead of the pruned system (3), (4)]. The QKF is derived under the assumption that the vector of estimation errors of filtered states and exogenous innovations $\Theta_{t,t+1} \equiv (\omega'_t - E_t\omega'_t, \varepsilon'_{t+1})'$ is Gaussian. Ivashchenko (2014) assumes a linear updating rule similar to (11): $\omega_{t+1,t+1} = \omega_{t+1,t} + \vartheta_t \cdot (y_{t+1} - y_{t+1,t})$, where ϑ_t is defined analogously to ϕ_t in (12); thus knowledge of $V^\omega_{t+1,t}$ (one-period ahead conditional variance of ω_{t+1}), is required for the QKF filter. (2) implies that $V^\omega_{t+1,t}$ depends on the conditional *fourth* moments of Θ_{t+1}. Under the assumed normality of Θ_{t+1}, it is easy to compute those fourth moments in closed-form, as functions of the conditional variance of Θ_{t+1}.

[9] Recall that the observable y_{t+1} is a *linear* function of Z_{t+1} [see (7)]; this may help to explain the good performance of the linear updating rule.

3 Monte Carlo Evidence

3.1 A Textbook RBC Model

The method is tested for a basic RBC model. Assume a representative infinitely-lived household whose date t expected lifetime utility V_t is given by $V_t = \left\{\frac{1}{1-\sigma}C_t^{1-\sigma} - \frac{1}{1+1/\eta}N_t^{1+1/\eta}\right\} + \lambda_t \beta E_t V_{t+1}$, where C_t and N_t are consumption and hours worked, at t, respectively. $\sigma > 0$ and $\eta > 0$ are the risk aversion coefficient and the (Frisch) labor supply elasticity. λ_t is an exogenous random taste (discount factor) shock that equals unity in steady state. $0 < \beta < 1$ is the steady state subjective discount factor. The household maximizes expected lifetime utility subject to the period t resource constraint

$$C_t + I_t = Y_t, \tag{13}$$

where Y_t and I_t are output and gross investment, respectively. The production function is

$$Y_t = \theta_t K_t^{\alpha} N_t^{1-\alpha} \tag{14}$$

where K_t is the beginning-of-period t capital stock, and $\theta_t > 0$ is exogenous total factor productivity (TFP). The law of motion of the capital stock is

$$K_{t+1} = (1-\delta)K_t + I_t. \tag{15}$$

$0 < \alpha, \delta < 1$ are the capital share and the capital depreciation rate, respectively. The household's first-order conditions are:

$$\lambda_t E_t \beta (C_{t+1}/C_t)^{-\sigma}(\theta_{t+1}\alpha K_{t+1}^{\alpha-1}N_{t+1}^{1-\alpha} + 1 - \delta) = 1, \quad C_t^{-\sigma}(1-\alpha)\theta_t K_t^{\alpha} N_t^{-\alpha} = N_t^{1/\eta}. \tag{16}$$

The forcing variables follow independent autoregressive processes:

$$\ln(\theta_t) = \rho_\theta \ln(\theta_{t-1}) + \varepsilon_{\theta,t}, \quad \ln(\lambda_t) = \rho_\lambda \ln(\lambda_{t-1}) + \varepsilon_{\lambda,t}, 0 < \rho_\theta, \rho_\lambda < 1, \tag{17}$$

where $\varepsilon_{\theta,t}$ and $\varepsilon_{\lambda,t}$ are normal i.i.d. white noises with standard deviations σ_θ and σ_λ.

The numerical simulations discussed below assume $\beta = 0.99, \eta = 4, \alpha = 0.3$, $\delta = 0.025$, $\rho_\theta = \rho_\lambda = 0.99$; parameter values in that range are standard in quarterly macro models. The risk aversion coefficient is set at a high value, $\sigma = 10$, so that the model has enough curvature to allow for non-negligible differences between the second-order accurate and linear model approximations. One model variant assumes shocks that are much larger than the shocks in standard macro models, in order to generate big differences between the two approximations: $\sigma_\theta = 0.20$, $\sigma_\lambda = 0.01$. I refer to this variant as the 'big shocks' variant. I also consider a second 'small shocks' variant, in which the standard deviations of shocks are twenty times smaller:

$\sigma_\theta = 0.01, \sigma_\lambda = 0.0005$ (conventional RBC models assume that the standard deviation of TFP innovations is about 1 %; e.g., Kollmann 1996).[10] The observables are assumed to be GDP, consumption, investment and hours worked; independent measurement error is added to log observables. Measurement error has a standard deviation of 0.04 (0.002) for each observable, in the model variant with big (small) shocks.

Chris Sims' MATLAB program gensys2 is used to compute first- and second-order accurate model solutions. The model is approximated in terms of logged variables. I apply the gensys2 algorithm to the 7-equations system (13)–(17) using the state vector $\omega_t \equiv (\ln(K_{t+1}), \ln(Y_t), \ln(C_t), \ln(I_t), \ln(N_t), \ln(\theta_t), \ln(\lambda_t))$. Simplifying the system (e.g. by plugging (14) and (15) into (13) and omitting Y_t, I_t) does not affect the results.

3.1.1 Predicted Standard Deviations

Table 1 reports unconditional standard deviations of 7 logged variables (GDP, consumption, investment, capital, hours, TFP and the taste shock λ) generated by the first- and second-order approximations.[11] Model variants with both shocks, and variants with just one type of shock, are considered; moments for non-HP (Hodrick-Prescott) filtered variables are shown, as well as moments of HP filtered variables (smoothing parameter: 1600).

In the 'big shocks' model variant, the standard deviations of endogenous variables are huge; e.g., with both shocks, the standard deviation of (non-HP filtered) GDP is 176 % (82 %) under the second-order (first-order) approximation; GDP is thus about twice as volatile under the second-order approximation (than in the linearized model).[12] The capital stock, investment and hours worked (non-HP filtered) are about one-half more volatile under the second-order approximation than under the linear approximation. By contrast, consumption volatility is similar across the two approximations. Consumption is much less volatile than GDP, due to the assumed high risk aversion of the household. The preference shock (λ) is the main source of fluctuations in the capital stock, GDP and investment; the TFP shock (θ) is the main driver of consumption. The correlation between the second- and first-order approximations of a given variable is noticeably below unity, in the model variant with big shocks: e.g., about 0.7 for capital and investment, and 0.5 for GDP.

The 'small shocks' model variant generates much smaller standard deviations of endogenous variables that are roughly in line with predicted moments reported in the RBC literature (e.g., Kollmann 1996); e.g., the predicted standard deviation of HP-filtered GDP and investment are about 1 and 5 %, respectively (with both

[10] The relative size of the TFP and taste shocks assumed here (i.e. σ_θ 20-times larger than σ_ψ) ensures that each shock accounts for a non-negligible share of the variance of the endogenous variables; see below.

[11] The statistics in Table 1 are shown for variables without measurement error. The ranking of volatilities generated by the two approximations and shocks is not affected by the presence of measurement error.

[12] HP filtered variables are markedly less volatile than non-HP filtered variables; however, volatility remains much higher under the second-order approximation than under the linear approximation, in the 'big shocks' variant. E.g. the standard dev. of HP filtered GDP is 47 % (23 %) under the second- (first-) order approximation.

Table 1 RBC model: predicted standard deviations

	Y	C	I	K	N	θ	λ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(a) Model variant with big shocks ($\sigma_\theta = 0.20, \sigma_\lambda = 0.01$)							
(a.1) ***Non-HP filtered variables***							
Second-order model approximation							
Both shocks	1.757	0.300	5.366	3.400	2.609	1.418	0.071
Just θ shock	0.492	0.264	1.387	0.959	1.761	1.418	0.000
Just λ shock	1.558	0.133	4.799	3.096	1.762	0.000	0.071
Linearized model							
Both shocks	0.817	0.276	3.269	2.364	1.862	1.418	0.071
Just θ shock	0.469	0.264	1.285	0.929	1.751	1.418	0.000
Just λ shock	0.669	0.083	3.006	2.174	0.634	0.000	0.071
(a.2) ***HP filtered variables***							
Second-order model approximation							
Both shocks	0.469	0.053	1.962	0.115	0.688	0.259	0.013
Just θ shock	0.124	0.037	0.483	0.038	0.212	0.259	0.000
Just λ shock	0.420	0.034	1.706	0.104	0.608	0.000	0.013
Linearized model							
Both shocks	0.229	0.041	1.059	0.095	0.350	0.259	0.013
Just θ shock	0.118	0.037	0.416	0.037	0.205	0.259	0.000
Just λ shock	0.196	0.016	0.974	0.087	0.284	0.000	0.013
(b) Model variant with small shocks ($\sigma_\theta = 0.01, \sigma_\lambda = 0.0005$)							
(b.1) ***Non-HP filtered variables***							
Second-order model approximation							
Both shocks	0.041	0.014	0.164	0.118	0.093	0.071	0.004
Just θ shock	0.023	0.013	0.064	0.046	0.088	0.071	0.000
Just λ shock	0.034	0.004	0.151	0.109	0.032	0.000	0.004
Linearized model							
Both shocks	0.041	0.014	0.163	0.118	0.093	0.071	0.004
Just θ shock	0.023	0.013	0.064	0.046	0.088	0.071	0.000
Just λ shock	0.033	0.004	0.150	0.109	0.032	0.000	0.004
(b.2) ***HP filtered variables***							
Second-order model approximation							
Both shocks	0.011	0.002	0.053	0.005	0.018	0.013	0.001
Just θ shock	0.006	0.002	0.021	0.002	0.010	0.013	0.000
Just λ shock	0.010	0.001	0.049	0.004	0.014	0.000	0.001
Linearized model							
Both shocks	0.011	0.002	0.053	0.005	0.018	0.013	0.001
Just θ shock	0.006	0.002	0.021	0.002	0.010	0.013	0.000
Just λ shock	0.010	0.001	0.049	0.004	0.014	0.000	0.001

Note: Standard deviations (SD) of logged variables (listed above Cols. (1)–(7)), without measurement error, are shown for the RBC model in Sect. 3.1. The SD were computed using the formulae in the Appendix. SD are *not* reported in %.

Panel (a) ('Big shocks') assumes SD of innovations to TFP θ and to the taste parameter λ of 20 and 1 %, respectively. Panel (b) ('Small shocks') sets these SD at 1 and 0.05 %, respectively. Rows labeled 'Both shocks'; 'Just θ shock'; and 'Just λ shock' show moments predicted with simultaneous two shocks; with just the TFP shock; and with just the taste shock, respectively. Panels (a.1) and (b.1) report moments of Non-HP filtered variables; Panels (a.2) and (b.2) pertain to HP filtered variables (smoothing parameter: 1600). Y: GDP; C: consumption; I: gross investment; K: capital stock; N: hours worked

shocks). With small shocks, it remains true that variables are more volatile in the second-order model than in the linearized model, however, the difference is barely noticeable. E.g., the ratio of the GDP [investment] standard dev. across the $2^{nd}/1^{st}$ order approximations is merely 1.005 [1.002].

3.1.2 Filter Accuracy

I generate 50 simulation runs of $T = 500$ and of $T = 100$ periods for the observables, using the second-order (pruned) state equation of the RBC model. Each run is initialized at the unconditional mean of the state vector. I apply the KalmanQ filter to the simulated observables (with measurement error). I also use Ivashchenko's (2014) Quadratic Kalman Filter 'QKF' and a conventional Kalman filter, referred to as 'KalmanL' (that treats the linearized model (3) as the true DGP). In addition, the standard particle filter (as described in An and Schorfheide 2007)—referred to a 'PF(p)', where p is the number of particles—is applied to the pruned state equation (4); for the simulation runs with T = 500 periods, 100,000 particles are employed; for runs with T = 100 periods, versions of the PF with 100,000 and with 500,000 particles are used.[13] Accuracy is evaluated for the 7 logged latent variables considered in Table 1.

In each simulation run $s = 1, \ldots, 50$, the root mean square error (RMSE) is computed, across all (logged) 7 variables, $RMSE_{s,All} \equiv (\frac{1}{7T}\sum_{i=1}^{7}\sum_{t=1}^{T}(\omega_{s,t}^{i} - \omega_{s,t,t}^{i})^2)^{1/2}$, and separately for each individual variable $i = 1, \ldots 7$, $RMSE_{s,i} \equiv (\frac{1}{T}\Sigma_{t=1}^{T}(\omega_{s,t}^{i} - \omega_{s,t,t}^{i})^2)^{1/2}$, where $\omega_{s,t}^{i}$ is the true date t value of variable i in run s, while $\omega_{s,t,t}^{i}$ is the filtered estimate (conditional expectation) of that variable, given the date t information set. Table 2 reports RMSEs that are averaged across simulation runs. In the Panels labeled 'Average RMSEs', Column (1) shows average RMSE, across all 7 variables, $\frac{1}{50}\sum_{s=1}^{50} RMSE_{s,ALL}$, while Cols. (2)–(8) separately show average RMSEs for each individual variable i, $\frac{1}{50}\sum_{s=1}^{50} RMSE_{s,i}$. Also reported are maximum estimation errors across all variables, periods and runs, as well as maximum estimation errors for each variable i (across all periods and all simulation runs); see Panels labeled 'Maximum errors'. These accuracy measures are reported for each of the filters (see rows labeled 'KalmanQ', 'QKF', 'PF(100,000)', 'PF(500,000)', and 'KalmanL'). In addition, I report the fraction of simulation runs in which the KalmanQ filter generates lower RMSEs and lower maximum estimation error than the other filters.

Table 2 shows that the KalmanQ filter is more accurate than the PF and KalmanL filters, in all (or almost all) simulation runs—this holds for both the 'big shocks' and 'small shocks' model variants.[14] In *all* 50 simulation runs for the 'big shocks' model

[13] I apply KalmanL to de-meaned series, as the linearized model implies that the unconditional mean of state variables, expressed as differences from steady state, is zero, while variables generated from the second-order model have a non-zero mean. The initial particles used for the particle filter are drawn from a multi-variate normal distribution whose mean and variance are set to unconditional moments of the state vector.

[14] I also computed median absolute errors (MAEs) for the filtered series. The results (available on request) confirm the greater accuracy of the KalmanQ filter.

Table 2 RBC model: accuracy of filters

	All variables	Y	C	I	K	N	θ	λ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(a) Model variant with big shocks ($\sigma_\theta = 0.20, \sigma_\lambda = 0.01$)								
(a.1) 50 simulation runs with T = 500 periods								
Average RMSEs								
KalmanQ	0.157	0.038	0.006	0.040	0.387	0.039	0.127	0.022
QKF	–	–	–	–	–	–	–	–
PF (100,000)	1.189	1.022	0.188	0.907	2.193	1.447	0.873	0.070
KalmanL	1.939	1.659	0.184	1.605	4.183	0.770	1.143	0.082
Maximum errors								
KalmanQ	3.448	0.209	0.024	0.164	3.448	0.166	1.198	0.236
QKF	–	–	–	–	–	–	–	–
PF (100,000)	20.188	15.750	1.927	10.451	20.188	19.842	4.236	0.563
KalmanL	14.824	14.824	1.275	7.849	13.071	3.820	12.039	0.568
Fraction of runs in which RMSE is lower for KalmanQ than for other filters								
QKF	–	–	–	–	–	–	–	–
PF (100,000)	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.98
KalmanL	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.96
Fraction of runs in which maximum error is lower for KalmanQ than for other filters								
QKF	–	–	–	–	–	–	–	–
PF (100,000)	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.98
KalmanL	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.94
(a.2) 50 simulation runs with T = 100 periods								
Average RMSEs								
KalmanQ	0.176	0.039	0.006	0.039	0.435	0.039	0.141	0.023
QKF	6.994	0.142	0.025	0.065	17.386	0.130	2.383	4.962
PF (100,000)	0.828	0.755	0.121	0.744	1.370	0.892	0.679	0.049
PF (500,000)	0.597	0.527	0.108	0.499	0.892	0.695	0.658	0.030
KalmanL	1.917	1.448	0.176	2.063	3.955	0.973	0.975	0.067
Maximum errors								
KalmanQ	2.855	0.160	0.023	0.134	2.855	0.163	0.901	0.226
QKF	1438.080	8.138	1.500	5.763	1438.080	5.482	153.500	606.064
PF (100,000)	13.298	10.615	1.364	9.944	13.190	13.298	2.846	0.414
PF (500,000)	8.018	4.548	0.560	8.018	5.680	4.577	3.425	0.235
KalmanL	11.866	9.826	0.655	8.399	11.866	3.484	5.093	0.332
Fraction of runs in which RMSE is lower for KalmanQ than for other filters								
QKF	1.00	1.00	1.00	0.75	1.00	0.89	1.00	1.00
PF (100,000)	1.00	1.00	1.00	1.00	0.84	1.00	1.00	0.64
PF (500,000)	1.00	1.00	1.00	1.00	0.84	1.00	1.00	0.68
KalmanL	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.98
Fraction of runs in which maximum error is lower for KalmanQ than for other filters								
QKF	1.00	1.00	0.93	0.64	1.00	0.68	1.00	1.00
PF (100,000)	0.96	1.00	1.00	1.00	0.84	1.00	1.00	0.80
PF (500,000)	1.00	1.00	1.00	1.00	0.80	1.00	1.00	0.76
KalmanL	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.96

Table 2 continued

	All variables	Y	C	I	K	N	θ	λ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(b) Model variant with small shocks ($\sigma_\theta = 0.01, \sigma_\lambda = 0.0005$)								
(b.1) ***50 simulation runs with T = 500 periods***								
Average RMSEs								
KalmanQ	0.0022	0.0007	0.0003	0.0019	0.0044	0.0019	0.0022	0.0001
QKF	0.0024	0.0007	0.0003	0.0019	0.0049	0.0019	0.0023	0.0001
PF (100,000)	0.0222	0.0068	0.0057	0.0109	0.0180	0.0405	0.0343	0.0007
KalmanL	0.0411	0.0163	0.0059	0.0583	0.0585	0.0397	0.0312	0.0018
Maximum errors								
KalmanQ	0.0475	0.0041	0.0011	0.0074	0.0475	0.0081	0.0163	0.0012
QKF	0.0601	0.0103	0.0011	0.0078	0.0601	0.0080	0.0216	0.0017
PF (100,000)	0.2437	0.0434	0.0330	0.0723	0.2438	0.2295	0.1939	0.0054
KalmanL	0.2293	0.0753	0.0248	0.2146	0.2294	0.1837	0.1512	0.0074
Fraction of runs in which RMSE is lower for KalmanQ than for other filters								
QKF	0.60	0.32	0.34	0.44	0.58	0.54	0.60	0.64
PF (100,000)	1.00	1.00	1.00	1.00	0.98	1.00	1.00	1.00
KalmanL	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Fraction of runs in which maximum error is lower for KalmanQ than for other filters								
QKF	0.54	0.60	0.50	0.52	0.54	0.46	0.64	0.56
PF (100,000)	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
KalmanL	0.98	1.00	1.00	0.98	0.92	1.00	0.98	0.94
(b.2) ***50 simulation runs with T = 100 periods***								
Average RMSEs								
KalmanQ	0.0042	0.0007	0.0003	0.0019	0.0099	0.0019	0.0035	0.0002
QKF	0.0049	0.0007	0.0003	0.0019	0.0121	0.0019	0.0041	0.0003
PF (100,000)	0.0244	0.0060	0.0051	0.0096	0.0368	0.0362	0.0313	0.0009
PF (500,000)	0.0223	0.0046	0.0039	0.0074	0.0398	0.0278	0.0271	0.0010
KalmanL	0.0508	0.0224	0.0078	0.0819	0.0716	0.0466	0.0357	0.0019
Maximum errors								
KalmanQ	0.0496	0.0033	0.0009	0.0073	0.0496	0.0068	0.0171	0.0012
QKF	0.0543	0.0048	0.0009	0.0069	0.0543	0.0069	0.0193	0.0015
PF (100,000)	0.1866	0.0319	0.0261	0.0551	0.1866	0.1848	0.1473	0.0046
PF (500,000)	0.2591	0.0224	0.0171	0.0402	0.2591	0.1189	0.1143	0.0065
KalmanL	0.3001	0.0918	0.0263	0.3001	0.2900	0.1846	0.1424	0.0083
Fraction of runs in which RMSE is lower for KalmanQ than for other filters								
QKF	0.66	0.12	0.26	0.46	0.66	0.48	0.66	0.66
PF (100,000)	1.00	1.00	1.00	1.00	0.94	1.00	1.00	0.98
PF (500,000)	1.00	1.00	1.00	1.00	0.90	1.00	1.00	0.92
KalmanL	1.00	1.00	1.00	1.00	0.94	1.00	1.00	0.90
Fraction of runs in which maximum error is lower for KalmanQ than for other filters								
QKF	0.64	0.24	0.46	0.60	0.64	0.48	0.56	0.66
PF (100,000)	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
PF (500,000)	1.00	1.00	1.00	1.00	0.92	1.00	1.00	0.96
KalmanL	0.98	1.00	1.00	1.00	0.94	1.00	1.00	0.90

Table 2 continued

Note: The Table reports average Root Mean Squared Errors (RMSEs) and maximum estimation errors of estimated latent states, across simulation runs (errors are *not* expressed in %), for different variants of the RBC model; accuracy is reported across all logged variables (Col. (1) 'All variables'), and separately for each logged variable (see Cols. (2)–(8) labeled 'Y', ..., 'λ'); see Sect. 3.1. for further information. KalmanQ: the filter for pruned second-order models developed in this paper; *QKF*: Quadratic Kalman Filter (Ivashchenko 2014) that assumes that the unpruned second-order model (2) is the true DGP; PF:(p) standard particle filter with p particles; *KalmanL*: standard Kalman filter that assumes that the linearized DSGE model is the true DGP.
Panel (a) ('Big shocks') assumes std. of TFP and taste shock innovations of 20 and 1 %, respectively; panel (b) ('Small shocks') sets these standard deviations at 1 and 0.05 %, respectively. Y: GDP; C: consumption; I: gross investment; K: capital stock; N: hours worked; θ: TFP; λ: taste shock (all variables are expressed in logs).
The time paths for filtered state estimates generated by the QKF exploded in all 50 simulation runs of the 'Big shocks' model variant with T = 500 periods; see Panel (a.1). Thus, no RMSEs and Maximum Errors are reported for the QKF in Panel (a.1). The QKF exploded in 44 % of the 50 simulation runs of the 'Big shocks' model variant with T = 100 periods; see Panel (a.2). The RMSEs and Maximum Errors for the QKF reported in Panel (a.2) pertain to the 56 % of the simulation runs in which the QKF did *not* explode; the fraction of runs in which RMSEs and Maximum Errors are lower for KalmanQ than for QKF only pertains to the runs in which QKF did *not* explode.
The QKF did not explode in the simulation runs of the 'Small shocks' model variant (Panel (b))

variant with T=500 periods (see Panel (a.1)), the QKF generated time paths of filtered state estimates that diverged explosively from the true states.[15] That problem never arose for the KalmanQ filter. In the 'big shocks' model variant with T=100 periods (see Panel (a.2)), the QKF exploded in 22 of the 50 simulation runs (44 %); for that variant, the average RMSEs and maximum errors reported for the QKF pertain to the simulation runs in which the QKF did not explode; in those simulation runs the KalmanQ filter is markedly more accurate than the QKF. Average RMSEs generated by KalmanQ are often orders of magnitudes smaller than the RMSEs generated by the particle filter, and that even when 500,000 particles are used. E.g., for the simulation runs of the 'big shocks' model variant with T=100 periods, the average RMSEs for GDP are 0.039, 0.142, 0.755, 0.527 and 1.488, respectively, for KalmanQ, QKF, PF(100,000), PF(500,000), and KalmanL; the corresponding maximum errors are 0.160, 8.138, 10.615, 4.548 and 9.826, respectively (see Panel (a.2), Col. (2)).

In the 'small shocks' simulation runs (Table 2, Panels (b.1) and (b.2)), all the filters are more accurate than in the 'big shocks' simulations, and thus the *absolute* accuracy differences between the filters are smaller. The QKF did not explode in the 'small shocks' simulations. The filtered estimates of latent states generated by the KalmanQ filter and by the QKF are now broadly similar (across all variables, the KalmanQ filter is more accurate than the QKF in 54–66 % of all simulation runs; see Column (1)). For the 'small shocks' simulation runs with T=100 periods, the average RMSEs for GDP are 0.0007, 0.0007, 0.0060, 0.0046

[15] Once the QKF filtered estimates of the second-order accurate state variables $\omega_{t,t}$ reach large values, the one-step ahead covariance matrix $V^{\omega}_{t+1,t}$ takes huge values too (in the QKF, $V^{\omega}_{t+1,t}$ depends on $\omega_{t,t}$); at that point the observations are no longer able to correct the filtered series, and the filtered series may start to diverge explosively.

and 0.0224, for KalmanQ, QKF, PF(100,000), PF(500,000) and KalmanL respectively, while corresponding maximum errors are 0.0033, 0.0048, 0.0319, 0.0224 and 0.0918 (see Panel (b.2)). The *relative* improvement in accuracy from using the KalmanQ filter thus remains sizable, relative to the particle filter, and relative to KalmanL.

Note that the accuracy checks considered so far are based on pruned simulated sample paths (generated using (3),(4)). It seems interesting to also apply the filters to sample paths generated by the *unpruned* model (2). For the 'big shocks' model variant, *all* unpruned sample paths of length T=100 and T=500 explode. By contrast, for the 'small shocks' model variant, the unpruned sample paths fail to explode—in fact, those paths are highly correlated with the pruned sample paths; the performance of the filters is hence similar to that reported in Table 2, for the 'small shocks' pruned model variant.[16] Detailed results are available on request.

3.1.3 Computing Time

KalmanQ, QKF, the particle filters with 100,000 and 500,000 particles and KalmanL require 0.03, 0.05, 14.69, 81.21 and 0.01 seconds, respectively, to filter simulated series of T=100 periods generated by the RBC model, on a desktop computer with a 64-bit operating system and a 3.4 Ghz processor. For series of T=500 periods, the corresponding computing times are 0.12, 0.21, 73.72, 401.58 and 0.04 seconds, respectively. Thus, the KalmanQ filter is about 500 (3000) times faster than the particle filter with 100,000 (500,000) particles, and approximately 40 % faster than the QKF.

For a sufficiently large number of particles, the particle filter is (asymptotically) an *exact* algorithm for computing the conditional expectation of the state vector. However, the experiments in Table 2 suggest that a very large number of particles (above 500,000) is needed to outperform KalmanQ; the computational cost of using such a large number of particles would be substantial.

3.1.4 Evaluating the Quasi-Maximum Likelihood (QML) Parameter Estimates

For 50 simulations runs of the 'big shocks' model variant and of the 'small shocks' variant, with T=100 periods, I computed QML estimates of the risk aversion coefficient (σ), the labor supply elasticity (η), the autocorrelations of the forcing variables (ρ_θ, ρ_λ) and the standard deviations of the innovations to the forcing variables (σ_θ, σ_λ). (As before, four observables are assumed: GDP, consumption, investment and hours.) Table 3 reports the mean and median parameter estimates, and the standard deviation of the parameter estimates, across the sample of 50 estimates per model variant. The parameters are tightly estimated; mean and median parameter estimates are close to the true parameter values.[17]

[16] For the unpruned 'small shocks' sample paths, the KalmanQ filter and the QKF thus have similar performance; with T=100 [T=500], the KalmanQ filter is more accurate (across all variables) than the QKF in about 45 % [54 %] of the simulation runs.

[17] A more detailed evaluation of the small sample properties of the QML estimator is left for future research.

Table 3 RBC model: distribution of quasi-maximum likelihood estimates of model parameters based on the KalmanQ filter

	σ	η	ρ_θ	ρ_λ	σ_θ	σ_λ
	(1)	(2)	(3)	(4)	(5)	(6)
(a) Model variant with big shocks ($\sigma_\theta = 0.20, \sigma_\lambda = 0.01$)						
True parameter value	10.00	4.00	0.99	0.99	20.00 %	1.00 %
Mean estimate	10.41	4.13	0.99	0.98	18.73 %	0.97 %
Median estimate	10.09	4.01	0.99	0.98	18.71 %	0.92 %
Standard dev. of estimates	1.36	1.53	0.003	0.005	1.44 %	0.21 %
(b) Model variant with small shocks ($\sigma_\theta = 0.01, \sigma_\lambda = 0.0005$)						
True parameter value	10.00	4.00	0.99	0.99	1.00 %	0.05 %
Mean estimate	10.26	4.85	0.99	0.69	0.76 %	0.20 %
Median estimate	10.15	3.40	0.99	0.69	0.74 %	0.18 %
Standard dev. of estimates	1.20	3.39	0.005	0.26	0.12 %	0.14 %

Note: The Table reports true parameter values, as well as the mean, median and standard deviation of quasi-maximum likelihood (QML) estimates of structural model parameters based on the KalmanQ filter (developed in this paper), obtained for 50 simulation runs with T = 100 periods of the 'big shocks' RBC model variant (Panel (a)) and for 50 simulation runs (T = 100) of the 'small shocks' RBC model variant (Panel (b)). Four observables (GDP, consumption, investment and hours) are used for the QML estimation. σ: risk aversion coefficient (Col. (1)); η: Frisch labor supply elasticity (Col. (2)); ρ_θ [ρ_λ]: autocorrelation of TFP [taste shock] (Cols. (3)–(4)); σ_θ [σ_λ]: standard deviation of TFP [taste shock] innovation (Cols. (5)–(6))

3.2 State Equations with Randomly Drawn Coefficients

Many other Monte Carlo experiments confirmed that the KalmanQ filter is competitive with the particle filter, in terms of accuracy of the estimated state variables. To document the performance of the filter in a broad range of setting, I applied it to simulated data generated using variants of the pruned state equations (3),(4) whose coefficients were drawn randomly from normal distributions. Tables 4 reports moments of the resulting simulated latent state variables, while Table 5 documents the accuracy of the filters. In both Tables, Panel (a) pertains to models with $n = 20$ variables, while Panel (b) assumes $n = 7$ variables; I refer to the models in Panels (a) and (b) as 'medium models', and as 'small models', respectively. In both set-ups, $m = 7$ independent exogenous innovations are assumed, and the first four elements of the state vector ω_t are observed with measurement error ($n_Y = 4$). The standard deviations of the (independent) exogenous innovations (ε_{t+1}) and of measurement errors (ψ_t) are set at 1 %.[18] The elements of F_0 are independent draws from $N(0, 1)$ that are scaled by a common factor so that the largest element of F_0 is $(0.01)^2$ in absolute value. The elements of F_1 are independent draws from $N(0, 1)$ that are scaled by a common factor so that the largest eigenvalue of F_1 has an absolute value of 0.99. The elements of F_2 are independent draws from $N(0, 1)$. In one set of simulations, referred to as 'strong cur-

[18] $\Sigma_\varepsilon = (.01)^2 I_m$; $\Sigma_\psi = (.01)^2 I_{n_y}$. As before, the parameter ξ that scales the size of the shocks is normalized as $\xi = 1$.

Table 4 Models with randomly drawn coefficients: average standard deviations of state variables

	Strong curvature	Weak curvature
	(1)	(2)
(a) Medium models (n= 20)		
(a.1) ***Non-HP filtered variables***		
Second-order approximated model	3.061	0.119
Linearized model	0.106	0.106
(a.2) ***HP filtered variables***		
Second-order approximated model	1.089	0.094
Linearized model	0.092	0.092
(b) Small models ($n = 7$)		
(b.1) ***Non-HP filtered variables***		
Second-order approximated model	2.446	0.147
Linearized model	0.144	0.144
(b.2) ***HP filtered variables***		
Second-order approximated model	0.678	0.116
Linearized model	0.115	0.115

Note: The Table reports average standard deviations of simulated state variables generated by pruned state equations (see (3), (4)) whose coefficients were drawn randomly from normal distributions. Panel (a) considers models with $n = 20$ variables ('medium models'), while Panel (b) pertains to models with $n = 7$ variables ('small models'). In Column (1), labeled 'Strong curvature', all elements of the matrices of curvature coefficients F_{11}, F_{12}, F_{22} are independent draws from N(0,1); in Column (2), labeled 'Weak curvature', F_{11}, F_{12}, F_{22} are independent draws from $N(0, (0.01)^2)$, so that curvature is much smaller, on average than in Column (1). (See Sect. 3.2 for further details.)
For both the 'medium' and 'small' model variants, 50 random 'strong curvature' coefficient sets, and 50 random 'weak curvature' coefficient sets were drawn. For each of the resulting 200 sets of coefficients, the model was simulated over T = 100 periods. For each simulation run, the standard deviation of each state variable (i.e. of each element of the vectors ω_t and $\omega_t^{(1)}$) was computed (without measurement error); then, standard deviations were averaged across all variables and coefficient draws, for each of the four model classes (medium/small models with strong/weak curvature), respectively. The Table shows those averaged standard deviations. (The averaged standard deviations are *not* reported in %.)
The rows labeled 'second-order approximated model' and 'linearized model' show average standard deviations of elements of the state vectors ω_t and of $\omega_t^{(1)}$, respectively, where ω_t is generated using the pruned quadratic state equation (4), while $\omega_t^{(1)}$ is generated using the linearized state equation (3)

vature' simulations, all elements of F_{11}, F_{12}, F_{22} are independent draws from $N(0, 1)$; in another set of simulations with 'weak curvature', the elements of F_{11}, F_{12}, F_{22} are independent draws from $N(0, (0.01)^2)$, so that curvature is much smaller, on average. For both the 'medium' and 'small' model variants, 50 random 'strong curvature' coefficient sets, and 50 random 'weak curvature' coefficient sets were drawn. Thus, 200 different random sets of coefficients (F_0, F_1, F_2, F_{11}, F_{12}, F_{22}) are considered. For each set of coefficients, the state equations (3) and (4) were simulated over T=100 periods (each run was initialized at the unconditional mean of the state vector), and the filters were applied to the observables (with measurement error).

Table 4 reports (averaged) standard deviations of the latent state variables for the 'medium' and 'small' model variants with 'strong curvature' (Col. (1)) and with 'weak

Table 5 Models with randomly drawn coefficients: accuracy of filters

	Strong curvature	Weak curvature
(a) Medium models ($n = 20$)		
Average RMSEs		
KalmanQ	0.155	0.0308
QKF	–	0.0320
PF (100,000)	31.243	0.0851
KalmanL	2.978	0.0523
Maximum errors		
KalmanQ	4.784	0.2789
QKF	–	0.3612
PF (100,000)	4440.00	9.3808
KalmanL	149.441	1.2315
Fraction of runs in which RMSE is lower for KalmanQ than for other filters		
QKF	–	0.38
PF (100,000)	1.00	0.98
KalmanL	1.00	0.94
Fraction of runs in which maximum error is lower for KalmanQ than for other filters		
QKF	–	0.52
PF (100,000)	1.00	0.76
KalmanL	0.98	0.74
(b) Small models ($n = 7$)		
Average RMSEs		
KalmanQ	0.035	0.0184
QKF	0.204	0.0185
PF (100,000)	38.082	0.0186
KalmanL	1.651	0.0409
Maximum errors		
KalmanQ	1.12	0.0943
QKF	41.93	0.1084
PF (100,000)	3263.28	0.0954
KalmanL	37.38	0.1265
Fraction of runs in which RMSE is lower for KalmanQ than for other filters		
QKF	1.00	0.42
PF (100,000)	0.76	0.76
KalmanL	1.00	0.96
Fraction of runs in which maximum error is lower for KalmanQ than for other filters		
QKF	1.00	0.52
PF (100,000)	0.82	0.50
KalmanL	1.00	0.66

Table 5 continued

Note: The Table reports average Root Mean Squared Errors (RMSEs) and maximum estimation errors of estimated latent variables produced by four filters, across simulation runs and state variables, for versions of the pruned state equation (4) with randomly drawn coefficients; see Sect. 3.2. Panel (a) ('medium models') assumes n = 20 variables; Panel (b) ('small model') assumes n = 7 variables. In 'strong curvature' ['weak curvature'] experiments (see Column 1 [2]), all elements of the matrices of curvature coefficients F_{11}, F_{12}, F_{22} are independent draws from $N(0, 1)[N(0, (0.01)^2)]$. (See Sect. 3.2 for further details.) For both the 'medium' and 'small' model variants, 50 random 'strong curvature' coefficient sets, and 50 random 'weak curvature' coefficient sets were drawn. For each of the resulting 200 sets of coefficient, the model was simulated over T = 100 periods.

For each simulation run, the RMSE and the maximal error was computed, for each of the 'n' estimated latent variables; then, RMSEs were averaged across variables and coefficient draws, for each of the four model classes (medium/small models with strong/weak curvature); maximum errors were likewise determined across all n variables, and across all draws, for each of the four model classes.

KalmanQ: the filter for pruned second-order models developed in this paper; *QKF*: Quadratic Kalman Filter (Ivashchenko 2014) that assumes that the unpruned second-order model (2) is the true DGP; PF(p): standard particle filter with $p$particles; *KalmanL*: standard Kalman filter that assumes that the linearized model is the true DGP.

The time paths for filtered estimates of the state variables generated by the QKF exploded in all 50 simulation runs of the 'medium models' with 'strong curvature'; thus, no RMSEs and Maximum Errors are reported for the Quadratic Kalman Filter (QKF) in Column (1) of Panel (a).

The QKF exploded in 48 % of the 50 simulation runs of the 'small models' with 'strong curvature'; the RMSEs and Maximum Errors for the QKF reported in Column (1) of Panel (b) pertain to the 52 % of the simulation runs in which the QKF did *not* explode; the reported fraction of runs in which RMSEs and Maximum Errors are lower for KalmanQ than for QKF likewise only pertains to the runs in which QKF did *not* explode.

The QKF did not explode in the simulation runs of the 'weak curvature' model variants (Column (2))

curvature' (Col. (2)). The rows labeled 'Second-order approximated model' and 'Linearized model' show (averaged) standard deviations of ω_t and of $\omega_t^{(1)}$, respectively, where ω_t was generated using the pruned quadratic state equation (4), while $\omega_t^{(1)}$ was generated using the linear state equation (3).[19] In 'strong curvature' model variants, the average predicted volatility of ω_t (second-order accurate) is several times larger than that of $\omega_t^{(1)}$. By contrast, in the 'weak curvature' variants, the volatility of the second-order accurate variables is only slightly higher than that of the first-order accurate variables.

Table 5 compares the accuracy of the KalmanQ, QKF, PF(100,000) and KalmanL filters, for each of the four model classes (medium/small models with strong/weak curvature). For each model class, the KalmanQ filter generates lower average RMSEs and lower maximum errors than the PF(100,000) and KalmanL filters. E.g., for the 'medium models', the average RMSEs of KalmanQ, PF(100,000) and KalmanL are 0.155, 31.243 and 2.978, respectively, under 'strong curvature', and 0.031, 0.085 and 0.052, respectively under 'weak curvature' (see Panel (a)).

The time paths of filtered estimates of state variables generated by the QKF exploded in all simulation runs for 'medium models' with 'strong curvature'. For 'small models' with 'strong curvature', the QKF exploded in 48 % of the simulation runs; for the

[19] For each simulation run, the standard deviation of each element of ω_t and $\omega_t^{(1)}$ was computed; then, standard deviations were averaged across variables and coefficient draws, for each of the four model classes (medium/small models with strong/weak curvature).

runs where the QKF did not explode, the KalmanQ filter is markedly more accurate than the QKF (see Column (1) of Panel (b), Table 5). In the 'weak curvature' model variants, by contrast, the QKF did not explode in any of the runs–KalmanQ is slightly more accurate than QKF in terms of average RMSEs and Maximum Errors across all simulation runs.[20]

KalmanQ, QKF, PF(100,000) and KalmanL require 0.10, 0.14, 24.78 and 0.02 s, respectively, to filter simulated series of T=100 periods generated by the 'small models'; for 'medium models' the corresponding computing times are 0.86, 0.17, 78.49 and 0.03 s, respectively. This confirms the finding that the KalmanQ filter is much faster than the particle filter. The KalmanQ filter is faster than the QKF for 'small models', but not for 'medium models'.

4 Conclusion

This paper has developed a novel approach for the estimation of latent state variables in DSGE models that are solved using a second-order accurate approximation and the 'pruning' scheme of Kim et al. (2008). By contrast to particle filters, no stochastic simulations are needed for the deterministic filter here; the present method is thus much faster than particle filters. The use of the pruning scheme distinguishes the filter developed here from Ivashchenko's (2014) deterministic Quadratic Kalman filter (QKF). The present filter performs well even in models with big shocks and high curvature. In Monte Carlo experiments, the filter developed here generates more accurate estimates of latent state variables than the standard particle filter, especially when the model has big shocks and high curvature. The present filter is also more accurate than a Kalman filter that treats the linearized model as the true DGP. Due to its high speed, the filter presented here is suited for the estimation of model *parameters*; a quasi-maximum likelihood procedure can be used for that purpose.

Acknowledgments I am very grateful to three anonymous referees for detailed and constructive comments. I also thank Martin Andreasen, Sergey Ivashchenko, Jinill Kim and Raf Wouters for useful discussions. Financial support from the National Bank of Belgium and from 'Action de recherche concertée' ARC-AUWB/2010-15/ULB-11 is gratefully acknowledged.

Appendix: Computing Moments of the State Vector (for KalmanQ Filter Formula)

The unconditional mean and variance of the state vector Z_{t+1} of the augmented state equation (5) are given by: $E(Z_{t+1}) = (I - G_1)^{-1}G_0$ and $V(Z_{t+1}) = G_1 V(Z_{t+1})G_1' + V(u_{t+1})$, respectively. Stationarity of Z_{t+1} (which holds under the assumption that all eigenvalues of F_1 are strictly inside the unit circle) implies

[20] I also examined filter performance using unpruned sample paths generated by the state equations with randomly drawn coefficients (results available in request). All 'strong curvature' model variants generated exploding unpruned sample paths. In the 'weak curvature' model variants, none of the unpruned sample paths exploded—pruned and unpruned sample paths were highly correlated; the QKF and KalmanQ filter showed similar performance.

$E(Z_{t+1}) = E(Z_t)$, $V(Z_{t+1}) = V(Z_t)$. Once $V(u_{t+1})$ has been determined, $V(Z_{t+1})$ can efficiently be computed using a doubling algorithm. Note that $\omega_t^{(1)} = \sum_{i=0}^{\infty} (F_1)^i F_2 \varepsilon_{t-i}$ and recall that

$$u_{t+1} \equiv G_2 \varepsilon_{t+1} + G_{12} \omega_t^{(1)} \otimes \varepsilon_{t+1} + G_{22}[P(\varepsilon_{t+1}) - E(P(\varepsilon_{t+1}))]. \tag{A.1}$$

$E(\omega_t^{(1)}) = 0, E(\omega_t^{(1)} \otimes \varepsilon_{t+1}) = 0, E((\omega_t^{(1)} \otimes \varepsilon_{t+1})\varepsilon'_{t+1}) = 0, E((\omega_t^{(1)} \otimes \varepsilon_{t+1}) P(\varepsilon_{t+1})') = 0$ hold as ε_{t+1} has mean zero and is serially independent. Hence, the covariances between the first and second right-hand side (rhs) terms in (A.1), and between the second and third rhs terms are zero. $\varepsilon_t \sim N(0, \Sigma_\varepsilon)$ implies that the unconditional mean of all third order products of elements of ε_{t+1} is zero (Isserlis' theorem): $E\varepsilon_{t+1}^i \varepsilon_{t+1}^j \varepsilon_t^k = 0$ for all $i, j, k = 1, \ldots, \mathrm{m}$, where ε_{t+1}^h is the h-th element of ε_{t+1}. Thus the covariance between the first and third rhs terms in (A.1) too is zero. Note that $V(\omega_t^{(1)} \otimes \varepsilon_{t+1}) = V(\omega_t^{(1)}) \otimes \Sigma_\varepsilon$, with $V(\omega_t^{(1)}) = F_1 V(\omega_t^{(1)}) F_1' + F_2 \Sigma_\varepsilon F_2'$. Thus,

$$V(u_{t+1}) = G_2 \Sigma_\varepsilon G_2' + G_{12}(V(\omega_t^{(1)}) \otimes \Sigma_\varepsilon) G_{12}' + G_{22} V(P(\varepsilon_{t+1})) G_{22}'.$$

$\varepsilon_t \sim N(0, \Sigma_\varepsilon)$ also implies that the covariance between $\varepsilon_{t+1}^i \varepsilon_{t+1}^j$ and $\varepsilon_{t+1}^r \varepsilon_{t+1}^s$ is

$$Cov(\varepsilon_{t+1}^i \varepsilon_{t+1}^j, \varepsilon_{t+1}^r \varepsilon_{t+1}^s) = \sigma_{i,r}\sigma_{j,s} + \sigma_{i,s}\sigma_{j,r} \text{for i,j,r,s} = 1, \ldots, \mathrm{m},$$

where $\sigma_{i,r} = E(\varepsilon_{t+1}^i \varepsilon_{t+1}^r)$. (See, e.g., Triantafyllopoulos 2002). This formula allows to compute $V(P(\varepsilon_{t+1}))$, the covariance matrix of the vector

$$P(\varepsilon_{t+1}) \equiv (\varepsilon_{t+1}^1 \varepsilon_{t+1}^1, \varepsilon_{t+1}^1 \varepsilon_{t+1}^2, \ldots, \varepsilon_{t+1}^1 \varepsilon_{t+1}^m, \varepsilon_{t+1}^2 \varepsilon_{t+1}^2, \ldots, \varepsilon_{t+1}^2 \varepsilon_{t+1}^m, \ldots, \varepsilon_{t+1}^{m-1} \varepsilon_{t+1}^{m-1}, \varepsilon_{t+1}^{m-1} \varepsilon_{t+1}^m, \varepsilon_{t+1}^m \varepsilon_{t+1}^m).$$

Conditional Variance of State-Space Disturbance

To derive the formula for the conditional variance of u_{t+1} [(10) in main text] these facts are used:

(i) $E((\omega_t^{(1)} \otimes \varepsilon_{t+1})\varepsilon'_{t+1} | \Upsilon^t) = \omega_{t,t}^{(1)} \otimes \Sigma_\varepsilon$, with $\omega_{t,t}^{(1)} \equiv E(\omega_t^{(1)} | \Upsilon^t)$.

(ii) $E((\omega_t^{(1)} \otimes \varepsilon_{t+1})(\omega_t^{(1)} \otimes \varepsilon_{t+1})' | \Upsilon^t) = E((\omega_t^{(1)} \omega_t^{(1)\prime}) \otimes (\varepsilon_{t+1}\varepsilon'_{t+1}) | \Upsilon^t) = E((\omega_t^{(1)} \omega_t^{(1)\prime}) | \Upsilon^t) \otimes \Sigma_\varepsilon = (V_{t,t}^{\omega^{(1)}} + \omega_{t,t}^{(1)} \omega_{t,t}^{(1)\prime}) \otimes \Sigma_\varepsilon$. (Note that $V_{t,t}^{\omega^{(1)}} = E(\omega_t^{(1)} \omega_t^{(1)\prime} | \Upsilon^t) - E(\omega_t^{(1)} | \Upsilon^t) E(\omega_t^{(1)} | \Upsilon^t)' = E(\omega_t^{(1)} \omega_t^{(1)\prime} | \Upsilon^t) - \omega_{t,t}^{(1)} \omega_{t,t}^{(1)\prime}$.)

(iii) $E(P(\varepsilon_{t+1})\varepsilon'_{t+1} | \Upsilon^t) = 0$, $E(P(\varepsilon_{t+1})(\omega_t^{(1)} \otimes \varepsilon_{t+1})' | \Upsilon^t) = 0$ (due to Isserlis' theorem). Thus, the conditional covariance between the 1st and 3rd rhs terms in (A.1) and between the 2nd and 3rd rhs terms is zero.

References

An, S., & Schorfheide, F. (2007). Bayesian analysis of DSGE models. *Econometric Reviews*, *26*, 113–172.

Andreasen, M. (2012). Non-linear DSGE models and the central difference Kalman filter. *Journal of Applied Econometrics*, *28*, 929–955.

Andreasen, M., Fernández-Villaverde, J., & Rubio-Ramírez, J. (2013). *The pruned state-space system for non-linear DSGE models: Theory and empirical applications*. NBER Working Paper No. 18983.

Aruoba, S., Bocola, L., & Schorfheide, F. (2012). *A new class of nonlinear time series models for the evaluation of DSGE models*. Working Paper, University of Maryland and University of Pennsylvania.

Collard, F., & Juillard, M. (2001). *Perturbation methods for rational expectations models*. Working Paper, GREMAQ and CEPREMAP.

Del Negro, M., & Schorfheide, F. (2011). Bayesian macroeconometrics. In J. Geweke, G. Koop, & H. van Dijk (Eds.), *The Oxford handbook of Bayesian econometrics* (pp. 293–389). Oxford: Oxford University Press.

Fernández-Villaverde, J., & Rubio-Ramírez, J. (2007). Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies*, *74*, 1059–1087.

Hamilton, J. (1994). *Time series analysis*. Princeton: Princeton University Press.

Harvey, A. (1989). *Forecasting-structural time series models and the Kalman filter*. Cambridge: Cambridge University.

Ivashchenko, S. (2014). DSGE model estimation on the basis of second-order approximation. *Computational Economics*, *43*, 71–82

Jin, H., & Judd, K. (2000). *Applying PertSolv to complete market RBC models*. Working paper, Stanford University.

Julier, S., & Uhlmann, J. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, *92*, 401–422.

Kim, J., Kim, S., Schaumburg, E. & Sims, C. (2008). Calculating and using second-order accurate solutions of discrete-time dynamic equilibrium models. *Journal of Economic Dynamics and Control*, *32*, 3397–3414.

Kollmann, R. (1996). Incomplete asset markets and the cross-country consumption correlation puzzle. *Journal of Economic Dynamics and Control*, *20*, 945–962.

Kollmann, R. (2002). Monetary policy rules in the open economy: Effects of welfare and business cycles. *Journal of Monetary Economics*, *49*, 989–1015.

Kollmann, R. (2004). *Solving non-linear rational expectations models: Approximations based on Taylor expansions*. Working Paper, University of Paris XII.

Kollmann, R., Kim, S., & Kim, J. (2011). Solving the multi-country real business cycle model using a perturbation method. *Journal of Economic Dynamics and Control*, *35*, 203–206.

Lombardo, G., & Sutherland, A. (2007). Computing second-order accurate solutions for rational expectation models using linear solution methods. *Journal of Economic Dynamics and Control*, *31*, 515–530.

Norgaard, M., Poulsen, N., & Ravn, O. (2000). New developments in state estimation for nonlinear systems. *Automatica*, *36*, 1627–1638.

Schmitt-Grohé, S., & Uribe, M. (2004). Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, *28*, 755–775.

Sims, C. (2000). *Second order accurate solution of discrete time dynamic equilibrium models*. Working Paper, Economics Department, Princeton University.

Smets, F., & Wouters, R. (2007). Shocks and Frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, *97*, 586–660.

Triantafyllopoulos, K. (2002). *Moments and cumulants of the multivariate real and complex Gaussian distributions*. Working Paper, University of Bristol.