# ON THE ACCURACY OF LINEAR DSGE SOLUTION METHODS AND THE CONSEQUENCES FOR LOG-NORMAL ASSET PRICING

ALEXANDER MEYER-GOHDE

ABSTRACT. This paper demonstrates an economically significant loss of accuracy in standard, generalized Schur (or QZ) decomposition based solutions methods for linear dynamic stochastic general equilibrium (DSGE) models. This is illustrated in a simple production-based asset pricing model with external habit formation, calibrated to match post-war US consumption growth and the equity premium. When there is insufficient eigenvalue separation about the unit circle, QZ-based numerical solutions miss the equity premium by up to several annualized percentage points. While none of the numerical solution methods gave any indication of this error, easily implementable backward-error metrics are shown to successfully warn of such potential inaccuracies.

*JEL classification codes*: C61, C63, E17

*Keywords*: Numerical accuracy; Production-based asset pricing; DSGE; Solution methods

## 1. INTRODUCTION

The asset pricing literature abounds with puzzles and perhaps the most prominent is the equity premium puzzle (Mehra and Prescott, 1985; Mehra, 2003) that seeks to understand how risky assets can command such a high excess return in the face of moderate underlying volatility. While many convincing consumption based explanations that modify assumptions on, say, the stochastic properties of the pricing kernel have been offered, production based asset pricing face the additional challenge of needing to provide a structural cause of these stochastic properties. Providing a structural explanation invariably requires solving a structural model, the most common being dynamic stochastic general equilibrium (DSGE) models, which generally need to be

solved numerically. Cochrane (2008, p. 300) expressed concern regarding the accuracy of solution approximations in general equilibrium and this paper points out a surprising degradation of the accuracy of solution approximations in the simplest approximation, linear approximations, and their consequences for the equity premium reported by these methods.

Providing a solution to a DSGE model involves solving a functional equation to determine an unknown function that maps the sequences of variables in the information set into the endogenous variables of the model, resolving expectations of these same endogenous variables (Judd, 1998; Fernández-Villaverde, Rubio-Ramírez, and Schorfheide, 2016). Linear DSGE models and associated linear solutions have long been studied, e.g., Blanchard (1979) and Blanchard and Kahn (1980), and modern numerical packages such as Dynare (Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot, 2011), Gensys (Sims, 2001), (Perturbation) AIM (Anderson and Moore, 1985; Anderson, Levin, and Swanson, 2006), Uhlig's Toolkit (Uhlig, 1999) and Solab (Klein, 2000) not only provide tools for solving a wide range of linear models, but also provide a first step in the solution procedure for many nonlinear methods as well.

The substantial hurdle in these linear methods is finding a solution to a (matrix) quadratic equation, frequently required to be the unique stable solution. For multivariate models with potentially singular coefficient matrices, the standard method is to double the dimension of the problem and employ the generalized Schur or QZ decomposition of Moler and Stewart (1973). While this algorithm is backward stable for the generalized eigenvalue decompositions for which it was designed, it is not always backward stable for quadratic eigenvalue problems (Tisseur, 2000) and may yield ill-conditioned eigenvalues for quadratic matrix polynomials (Higham, Mackey, and Tisseur, 2006; Higham, Mackey, Tisseur, and Garvey, 2008). I present the backward-forward error analysis of Higham and Kim (2001) for matrix quadratic equations to provide an assessment of the accuracy of various solution methods in the literature valid when a symbolic solution is not available for comparison.

I demonstrate potential numerical inaccuracy in a simple, two equation canonical real business cycle model with habit formation, chosen as its small dimension enables a symbolic solution to serve as a benchmark. Exploring calibrations that match two calibration targets, post-war US consumption growth volatility and the average equity premium, I find that standard QZ based methods deliver an equity premium often off by

several annual percentage points. Using the pseudospectrum (see Tisseur and Higham, 2001; Higham and Tisseur, 2002), the set of eigenvalues subjected to small perturbations, I demonstrate the numerical inaccuracies gain economic significance when the separation of stable and unstable eigenvalues at the unit circle becomes numerically small. Backward error diagnostics that can be calculated at minimal additional cost and in the absence of a symbolic or analytic solution successfully warn of potential inaccuracies. This is of immediate, practical use, as none of the algorithms from the literature I explore produced any warning that their solutions might suffer from economically significant losses of accuracy. More general production based explanations, such as Jermann (1998), Tallarini (2000), and Croce (2014), combine more involved specifications of preferences and mappings from exogenous variables to underlying macroeconomic variables, and must also invariably solve their models numerically. Inasmuch as the solutions thereto are potentially subject to this inaccuracy in the solution of the underlying DSGE models, so too are their asset pricing predictions.

The remainder of the paper is structured as follows. Section 2 introduces the real business cycle model and the log-linear asset pricing approach. In section 3 I turn to solution methods of linear DSGE models and in section 4 to the numerical mathematics literature on solving matrix quadratic and generalized eigenvalue problems to assess the accuracy of these solutions. Section 5 presents the results from various methods in the DSGE linear solution literature for several calibrations. Finally, I conclude in section 6.

## 2. A SIMPLE LOG NORMAL DSGE ASSET PRICING MODEL

Here I layout a simple production-based asset pricing model, based on a standard real business cycle model (Kydland and Prescott, 1982; King and Rebelo, 1999) with external habit formation and a power utility kernel. (Constantinides, 1990; Campbell and Cochrane, 1999; Campbell, 2003) The representative household seeks to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, X_t), \ 0 < \beta < 1 \tag{1}$$

where $c_t$ is consumption and $X_t$ the externl habit stock, subject to

$$c_t + k_t = e^{z_t} k_{t-1}^{\alpha} + (1-\delta) k_{t-1}, \ 0 < \alpha, \delta < 1 \tag{2}$$

where $k_t$ is the capital stock accumulated at time $t$ and $z_t$ is total factor productivity that follows the AR(1) process

$$z_t = \rho z_{t-1} + \omega \varepsilon_t, \ \varepsilon_t \overset{\text{i.i.d.}}{\sim} N(0,1), \ |\rho| < 1, \ 0 < \omega \tag{3}$$

The first order condition of the maximization problem is

$$1 = E_t \left[ \underbrace{\beta \frac{u_c(c_{t+1}, X_{t+1})}{u_c(c_t, X_t)}}_{m_{t+1}} \underbrace{\left( \alpha e^{z_{t+1}} k_t^\alpha + 1 - \delta \right)}_{R_{t+1}} \right] \tag{4}$$

where $m_{t+1}$ is the stochastic discount factor or pricing kernel and $R_{+1}$ is the (risky) return on capital. Assuming an external habit such that $X_t = c_{t-1}$ in equilibrium with $h$ the degree of habit formation and power or CRRA utility with risk coefficient $\sigma$, marginal utility is $u_c(c_t, X_t) = (c_t - h c_{t-1})^{-\sigma}$. Equations (2)-(4) characterize a equilibrium for the stochastic sequences $\{c_t, k_t, z_t\}_{t=0}^\infty$ given a sequence of shocks $\{\varepsilon_t\}_{t=0}^\infty$ and initial conditions $c_{-1}, k_{-1}, z_{-1}$.

Defining the steady state, values $\overline{c}, \overline{k}, \overline{z}$ that solve (2)-(4) with $\varepsilon_t = 0 \, \forall t$, equations (2) and (4) can be log-linearized around these values to yield

$$0 = A E_t[y_{t+1}] + B y_t + C y_{t-1} + D z_t, \ y_t = \begin{bmatrix} \hat{c}_t & \hat{k}_t \end{bmatrix}' \tag{5}$$

$$z_t = \rho z_{t-1} + \omega \varepsilon_t, \ \varepsilon_t \overset{\text{i.i.d.}}{\sim} N(0,1) \tag{6}$$

a 2 by 2 system of equations linear in the log-deviations of the endogenous variables, $c_t$ and $k_t$, from their steady states, $\hat{w}_t \equiv \log w_t - \log \overline{w}$, for $w \in c, k$.

Following Hansen and Singleton (1983); Campbell and Shiller (1988); Campbell (2003), risky (say, $R_t$ from above) and risk-free (via no arbitrage, $1 = E_t[m_{t+1}] R_t^f$) assets can be priced under the implied joint log-normality of the approximation above via

$$1 = E_t \left[ e^{\widehat{m_{t+1}} + \widehat{R_{t+1}}} \right], \ \text{and} \ 1 = \overline{R^f} e^{R_t^f} E_t \left[ e^{\widehat{m_{t+1}}} \right] \tag{7}$$

which gives the risk premium as $-cov_t(\widehat{m_{t+1}}, \widehat{R_{t+1}})$, following, e.g., Lettau (2003), and can be expressed in terms of the variance of $z_t$ ($\omega^2$) as $\frac{\sigma}{1-h} \alpha Q_{cz} \left( 1 + \beta(1-\delta) \right) \omega^2$. Importantly, the coefficient $Q_{cz}$, the impact of technology on (log) consumption, must be solved for numerically even in this (log) linear case. To this solution I turn to in the following section.

## 3. SOLUTION METHODS

Standard numerical solution packages available to economists and policy makers—e.g., Dynare (Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot, 2011), Gensys (Sims, 2001), (Perturbation) AIM (Anderson and Moore, 1985; Anderson, Levin,

and Swanson, 2006), Uhlig's Toolkit (Uhlig, 1999) and Solab (Klein, 2000)—all analyze models that in some way or another can be expressed in the form of the nonlinear functional equation

$$0 = E_t[f(y_{t+1}, y_t, y_{t-1}, \varepsilon_t)] \tag{8}$$

The model equations (optimality conditions, resource constraints, market clearing conditions, etc.) are represented by the $n_y$-dimensional vector-valued function $f : \mathbb{R}^{n_y} \times \mathbb{R}^{n_y} \times \mathbb{R}^{n_y} \times \mathbb{R}^{n_e} \to \mathbb{R}^{n_y}$; $y_t \in \mathbb{R}^{n_y}$ is the vector of $n_y$ endogenous variables; and $\varepsilon_t \in \mathbb{R}^{n_e}$ the vector of $n_e$ exogenous shocks with a known distribution, where $n_y$ and $n_e$ are positive integers ($n_y, n_e \in \mathbb{N}$).

The solution to (8) is sought as the unknown function

$$y_t = y(y_{t-1}, \varepsilon_t), \quad y : \mathbb{R}^{n_y + n_e} \to \mathbb{R}^{n_y} \tag{9}$$

a function in the time domain that maps states, $y_{t-1}$ and $\varepsilon_t$, into endogenous variables, $y_t$. An analytic form for (9) is rarely available and researchers and practitioners are compelled to find approximative solutions. However, a steady state, $\overline{y} \in \mathbb{R}^{n_y}$ a vector such $\overline{y} = y(\overline{y}, 0)$ and $0 = f(\overline{y}, \overline{y}, \overline{y}, 0)$ can frequently be recovered, either analytically or numerically, providing a point of expansion around which local solutions may be recovered.

A first-order, or linear approximation, of (8) at the steady state delivers, analogously to the log linearized model of section 2,

$$0 = A E_t[y_{t+1}] + B y_t + C y_{t-1} + D \varepsilon_t \tag{10}$$

where $A$, $B$, $C$, and $D$ are the derivatives of $f$ in (8) with respect to its arguments and, recycling notation, the $y$'s in (10) refer to (log) deviations of the endogenous variables from their steady states, $\overline{y}$.

In analogy to (9), the standard approach to finding a solution to the linearized model (10) is to find a linear solution in the form

$$y_t = P y_{t-1} + Q \varepsilon_t \tag{11}$$

a recursive solution in the time domain–solutions that posit $y_t$ as a function of its own past, $y_{t-1}$, and exogenous innovations, $\varepsilon_t$.

Inserting (11) into (10) and taking expectations ($E_t[\varepsilon_{t+1}] = 0$), yields the restrictions

$$0 = AP^2 + BP + C, \quad 0 = (AP + B)Q + D \tag{12}$$

Generally, a unique $P$ with eigenvalues inside the closed unit circle is sought (I will address this formally below). Lan and Meyer-Gohde (2014) prove the latter can be uniquely solved for $Q$ if such a $P$ can be found. Hence, the hurdle is the former, matrix quadratic equation.

To assist in the analysis, I will formalize the matrix quadratic equation in (12). For $A$, $B$, and $C \in \mathbb{R}^{n_y \times n_y}$, a matrix quadratic $M(P) : \mathbb{C}^{n_y \times n_y} \to \mathbb{C}^{n_y \times n_y}$ is defined as

$$M(P) \equiv A P^2 + B P + C \tag{13}$$

with its solutions, called solvents,[1] given by $P \in \mathbb{C}^{n_y \times n_y}$ if and only if $M(P) = 0$. The eigenvalues of the solvent, called latent roots of the associated lambda matrix[2] $M(\lambda) : \mathbb{C} \to \mathbb{C}^{n \times n}$ (here of degree two), are given via

$$M(\lambda) \equiv A \lambda^2 + B \lambda + C \tag{14}$$

The latent roots are (i) values of $\lambda \in \mathbb{C}$ such that $\det M(\lambda) = 0$ and (ii) $n_y - \mathrm{rank}(A)$ infinite roots. An explicit link between the quadratic matrix problem and the quadratic eigenvalue problem is given via

$$\lambda \in \mathbb{C} : \left( A \lambda^2 + B \lambda + C \right) x = 0 \text{ for some } x \neq 0 \tag{15}$$

which has been reviewed extensively by Tisseur and Meerbergen (2001) and for which Hammarling, Munro, and Tisseur (2013) provide a comprehensive method to improve the accuracy of its solutions. If a unique stable solution is sought or required, this can be formulated via an adaptation of Blanchard and Kahn's (1980) rank and order conditions to the matrix quadratic formulation above. First assume there exist $2n_y$ latent roots of (14) of which $n_y$ lie inside and $n_y$ outside the unit circle. Second, there exists an $P \in \mathbb{R}^{n_y \times n_y}$ such that $M(P) = 0$ and $|eig(P)| < 1$.

Most linear DSGE methods use a generalized Schur or QZ decomposition (Moler and Stewart, 1973; Golub and van Loan, 2013) of the companion linearization of (10)[3] in some form or another. For the formulation above, the matrix quadratic (12) can be brought into

---

[1]The analysis proceeds in the complex plane, but the results carry over when solutions are restricted to be real valued due to the eigenvalue separation about the unit circle assumed below, see also Klein (2000).

[2]See, e.g., Dennis, Jr., Traub, and Weber (1976, p. 835) or Gantmacher (1959, vol. I, p. 228).

[3]Instead of the method of undetermined coefficients taken for expediency here, a multivariate pivoted Blanchard (1979) approach that delivers the solution constructively is presented in the appendix.

the QZ form as

$$F \begin{bmatrix} I_{n_y} \\ P \end{bmatrix} P = G \begin{bmatrix} I_{n_y} \\ P \end{bmatrix}, \quad F \equiv \begin{bmatrix} I_{n_y} & 0_{n_y \times n_y} \\ 0_{n_y \times n_y} & A \end{bmatrix}, G \equiv \begin{bmatrix} 0_{n_y \times n_y} & I_{n_y} \\ -C & -B \end{bmatrix} \quad (16)$$

where $I_{n_y}$ is an $n_y \times n_y$ identity matrix and $0_{n_y \times n_y}$ is an $n_y \times n_y$ zero matrix.

Applying the QZ or generalized Schur decomposition (unitary $Q$ and $Z$ and upper triangular $S$ and $T$ with $Q^* F Z = S$ and $Q^* G Z = T$), Higham and Kim (Theorem 3 2000) prove that all solvents or solutions of (16) are of the form $P = Z_{21} Z_{11}^{-1} = Q_{11} S_{11}^{-1} T_{11} Q_{11}^{-1}$. The decomposition is intricately related to the quadratic eigenvalue problem (15) via

$$\lambda \in \mathbb{C} : (F\lambda - G) y, \text{ where } y = \begin{bmatrix} x' & x'\lambda \end{bmatrix} \text{ for some } x \neq 0 \quad (17)$$

$$\lambda \in \mathbb{C} : Q (S\lambda - T) \tilde{y}, \text{ where } \tilde{y} = Z^* \begin{bmatrix} x' & x'\lambda \end{bmatrix} \text{ for some } x \neq 0 \quad (18)$$

where the eigenvalues in both lines are identical following from unitary equivalence (Moler and Stewart, 1973) and hence identical to the eigenvalues in (15) and the latent roots of (14). From the upper triangularity of $S$ and $T$ it follows that the spectrum or set of eigenvalues of the pencil $P_{FG}(\lambda) = F\lambda - G$ is determined by the diagonal entries of $S$ and $T$

$$\rho(P_{FG}) = \{t_{ii}/s_{ii}, \text{ if } s_{ii} \neq 0; \infty, \text{ if } s_{ii} = 0; \emptyset, \text{ if } s_{ii} = t_{ii} = 0; i = 1,\ldots,2n_y\} \quad (19)$$

where $s_{ii}$ and $t_{ii}$ denote the $i$'th row and $i$'th column of $S$ and $T$ respectively.

Ordering the decomposition so that the eigenvalues outside the unit circle are in the lower right blocks of $S$ and $T$ (hence $S_{22}$ and $T_{22}$), the necessary and sufficient assumptions for a unique stable solution for $y_t$ of (10) to exist are (1) Regularity: $P_{FG}(z)$ is called regular if there exists a $z \in \mathbb{C}$ such that $det(Fz - G) \neq 0$; (2) Order: Of the $2n_y$ generalized eigenvalues, there are exactly $n_y$ stable roots inside the unit circle, and consequently, exactly $n_y$ unstable roots outside the unit circle; (3) Rank: $Z_{11}$, the upper right block of $Z$, is nonsingular. If and only if these three assumptions are fulfilled does a unique solution $P$ stable with respect to the closed unit circle exist. Consequentially, the overwhelming majority of the linear solution methods provided to researchers and practitioners in the standard numerical solution packages enumerated at the beginning of the section can be summarized by this single matrix decomposition.

Binder and Pesaran (1997), the cyclic reduction method in Dynare (Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot, 2011), and Anderson (2010) are three

prominent methods that solve for $P$ without appealing to the generalized Schur decomposition. Binder and Pesaran's (1997) "fully recursive method" works directly with the matrix quadratic (12) and iterates on

$$\tilde{P}_k = I_{n_y} - \tilde{A}\tilde{P}_{k-1}^{-1}\tilde{C}, \ \ \text{where } \tilde{A} \equiv B^{-1}A, \ \tilde{C} \equiv B^{-1}C, \ \tilde{P}_0 \equiv I_{n_y} \tag{20}$$

Delivering the solution to the matrix quadratic (12) as $P = -\tilde{P}_N^{-1}\tilde{C}$ for some maximum iteration $N$. The cyclic reduction method implemented in Dynare (Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot, 2011) operates on the following recursion (see Bini, Latouche, and Meini, 2002)

$$P = -\hat{A}_i^{-1}A_0 \tag{21}$$

where

$$\hat{A}_i = \hat{A}_{i-1} - A_{2,i-1}A_{1,i-1}^{-1}A_{0,i-1} \tag{22}$$

$$A_{1,i} = A_{1,i-1} - A_{0,i-1}A_{1,i-1}^{-1}A_{2,i-1} - A_{2,i-1}A_{1,i-1}^{-1}A_{0,i-1} \tag{23}$$

$$A_{0,i} = -A_{0,i-1}A_{1,i-1}^{-1}A_{0,i-1} \tag{24}$$

$$A_{2,i} = -A_{2,i-1}A_{1,i-1}^{-1}A_{2,i-1} \tag{25}$$

with initial conditions $\hat{A}_0 = B$, $A_{2,0} = A$, $A_{1,0} = B$, and $A_{0,0} = C$ until convergence of $\hat{A}_i$. Anderson (2010) applies the bi-orthogonality from the separation of stable and unstable solutions to solve for the left invariant space associated with unstable solutions via[4]

$$\begin{bmatrix} y_t \\ E_t[y_{t+1}] \end{bmatrix} = \begin{bmatrix} 0_{n_y \times n_y} & I_{n_y} \\ -A^{-1}C & -A^{-1}B \end{bmatrix} \begin{bmatrix} y_{t-1} \\ y_t \end{bmatrix} \Rightarrow \begin{bmatrix} V_1 & V_2 \end{bmatrix} \begin{bmatrix} 0_{n_y \times n_y} & I_{n_y} \\ -A^{-1}C & -A^{-1}B \end{bmatrix} = \mathcal{M} \begin{bmatrix} V_1 & V_2 \end{bmatrix} \tag{26}$$

where the vectors of $V$ span the invariant space associated with unstable eigenvalues. This gives $y_t = -V_2^{-1}V_1 y_{t-1}$ as the solution to the homogenous problem, i.e., the matrix quadratic (12), $P = -V_2^{-1}V_1$. Essentially, by rearranging or shuffling the equations and variables, Anderson (2010) is able to reformulate a potentially singular system requiring the generalized Schur decomposition into a nonsingular system that can be solved using standard eigenvalue methods. The key commonality of these three methods is that they avoid the QZ or generalized Schur decompositon.

---

[4]This assumes that $A$ is invertible, the general case can be found in Anderson (2010) and is merely slightly more involved, utilizing the shuffle-algorithm of Luenberger (1978) to yield an invertible $A$.

## 4. ACCURACY OF METHODS

In this section, I present different perspectives on the accuracy of the different methods presented in the previous section. As the model in section 2 is of low dimension, a symbolic solution is available. This allows me to directly compare the numerical solutions of the previous section with this symbolic solution. As such a symbolic solution is not available for larger models, this comparison is not useful in general. I accordingly present two further approaches. Firstly, I turn to the analysis of the pseudospectrum, which captures the numerical inaccuracy in the calculation of the spectrum or set of eigenvalues of a pencil, the foundation of the QZ method behind most of the methods of the previous section. Finally, I summarize a forward-backward error assessment of the matrix quadratic problem.

Apart from Anderson (2008), very little attention has been paid to comparing the accuracy of different algorithms for linear models[5] and to numerically addressing the assumptions necessary for the existence of a unique stable solution.[6] Improvements in the accuracy of the solution to linear DSGE models has implications for many nonlinear solutions as well. Anderson, Levin, and Swanson (2006) demonstrate that even small inaccuracies in lower orders compound to larger errors in the computation of higher, nonlinear solutions such as in Jin and Judd (2002).

Higham and Kim (2000) note that Matlab can reliably provide the solvents of matrix quadratic problems of a small dimensionality like the model of section 2, allowing for a direct comparison of the numerical solutions with this symbolic solution. Specifically, I will assess the different methods' solutions with respect to the symbolic solution via the two moments I will target in calibrating the model $E[rp]$ and $std(\log c_t)$ as well as the largest absolute deviations in the matrices for the linear solution or policy function (11), $P$ and $Q$,

$$\max(|\Delta P|) \equiv \max(|P_{symbolic} - P_{method}|), \ \ \max(|\Delta Q|) \equiv \max(|Q_{symbolic} - Q_{method}|) \quad (27)$$

---

[5]This is in stark contrast to the many studies that examine the accuracy of different nonlinear methods. See Fernández-Villaverde, Rubio-Ramírez, and Schorfheide (2016) for an overview.

[6]Heilberger, Klarl, and Maußner (2015) provides an exception, showing that, theoretically, if the rank assumption for the QZ decomposition is fulfilled for one ordering of the eigenvalues that conforms to the unit circle separation, it holds for any ordering that conforms to the same.

While useful as direct measures of accuracy, a symbolic solution for the policy function (11) is generally not available, even in this linear case. I will now turn to numerical methods to assess the accuracy of the solution.

The generalized Schur or QZ based methods from section 3 are eigenvalue based methods, with the triangular structure of the factorizations revealing the eigenvalues of the underlying inflated matrix pencil. Studies concerning the numerical robustness of generalized eigenvalue problems date back at least to Stewart (1972) and Wilkinson (1979), who provided examples of essentially arbitrary results from the QZ algorithm in the presence of nearly singular pencils, i.e. violation of the regularity assumption above. The computation of eigenvalues numerically is likewise subject to finite precision. Hammarling, Munro, and Tisseur (2013) provide a comprehensive study on improving the accuracy of quadratic eigenvalue problems. Anguas, Bueno, and Dopico (2019) provides a comparison of different conditioning numbers for the eigenvalues of matrix polynomials and conditioning numbers of polynomial eiqenvalues can be obtained via eigenvalues for perturbations of the polynomial or pseudospectra (see Tisseur and Higham, 2001; Higham and Tisseur, 2002). Specifically, Tisseur and Higham (2001), Mengi and Overton (2005), and Michiels, Green, Wagenknecht, and Niculescu (2006) apply pseudospectra to stability radii in continuous-time applications. Along with a comparison of the numerical and symbolic eigenvalues, $\lambda$,

$$\max(|\Delta eig|) \equiv \max(|\lambda_{symbolic} - \lambda_{method}|) \tag{28}$$

I will turn to the pseudospectrum for insight into the saddle-path stability vis-a-vis the unit circle in the problem laid out above.

Specifically, the pseudospectrum provides a perturbed analog to the spectrum or set of eigenvalues/latent roots of (14) and (15)

$$\rho_\epsilon(M) = \{\lambda \in \mathbb{C} : (M(\lambda) + \Delta M(\lambda))x = 0 \text{ for some } x \neq 0 \text{ and } \Delta M(\lambda) \tag{29}$$

$$\text{with } \|\Delta A\| \leq \epsilon\alpha_A, \|\Delta B\| \leq \epsilon\alpha_B, \|\Delta C\| \leq \epsilon\alpha_C\} \tag{30}$$

where $\Delta M(\lambda)$ represents the perturbation of the quadratic[7]

$$\Delta M(\lambda) \equiv \Delta A\,\lambda^2 + \Delta B\,\lambda + \Delta C \tag{31}$$

---

[7]This is perhaps easier to see via the identity $M(\lambda) + \Delta M(\lambda) = (A + \Delta A)\lambda^2 + (B + \Delta B)\lambda + (C + \Delta C)$.

and the $\alpha_i$'s control the perturbation, which are set as $\alpha_X = |X|$ using the 2-norm following Tisseur (2000). As shown in Tisseur and Higham (2001), this 2-norm definition of the pseudopectrum corresponds to the backward errors of the eigenvalues.

As proven in Tisseur (2000), while the QZ or generalized Schur algorithm is numerically stable for the generalized eigenvalue problem (Stewart, 1972), this is not the case for the quadratic eigenvalue problem, as it does not respect the structure of the latter. To see this, first define the pseudospectrum of (17) analogous to above

$$\rho_\epsilon(P_{FG}) = \{\lambda \in \mathbb{C} : (P_{FG}(\lambda) + \Delta P_{FG}(\lambda))x = 0 \text{ for some } x \neq 0 \text{ and } \Delta P_{FG}(\lambda) \tag{32}$$

$$\text{with } \|\Delta F\| \leq \epsilon \alpha_F, \|\Delta G\| \leq \epsilon \alpha_G\} \tag{33}$$

comparing the perturbations involved in (32) with (29)

$$\Delta P_{FG}(\lambda) \equiv \Delta F \lambda - \Delta G = \begin{bmatrix} \Delta F_{11} & \Delta F_{12} \\ \Delta F_{21} & \Delta F_{22} \end{bmatrix} \lambda - \begin{bmatrix} \Delta G_{11} & \Delta G_{12} \\ \Delta G_{21} & \Delta G_{22} \end{bmatrix} \tag{34}$$

$$\neq \left( \begin{bmatrix} I_{n_y} & 0_{n_y \times n_y} \\ 0_{n_y \times n_y} & \Delta A \end{bmatrix} \lambda - \begin{bmatrix} 0_{n_y \times n_y} & I_{n_y} \\ -\Delta C & -\Delta B \end{bmatrix} \right) \begin{bmatrix} I_{n_y} \\ I_{n_y} \lambda \end{bmatrix} = \begin{bmatrix} 0 \\ \Delta M(\lambda) \end{bmatrix} \tag{35}$$

Inspection underscores that, in general, perturbations of the QZ or generalized Schur of the companion linearization (16) do not respect the specific structure in the underlying matrix quadratic problem (12).

Higham and Kim (2001) provide bounds on the backward error and a condition number for the solvent, $P$, of a quadratic matrix equation. Beginning with the perturbed matrix quadratic

$$M(P + \Delta P) + \Delta M(P + \Delta P) \equiv (A + \Delta A)(P + \Delta P)^2 + (B + \Delta B)(P + \Delta P) + C + \Delta C \tag{36}$$

with normwise perturbations $\epsilon = \left\| \begin{bmatrix} \Delta A / \|A\|_F & \Delta B / \|B\|_F & \Delta C / \|C\|_F \end{bmatrix} \right\|_F$, they derive the sharp bound

$$\|\Delta P\|_F / \|P\|_F \leq \Psi(P)\epsilon + \mathcal{O}\left(\epsilon^2\right) \tag{37}$$

where $\Psi(P)$ is the condition number given by

$$\Psi(P) = \left\| \left(I_n \otimes AP + P' \otimes A + I_n \otimes B\right)^{-1} \left[ \|A\|_F \left(P^2\right)' \otimes I_n \quad \|B\|_F P' \otimes I_n \quad \|C\|_F I_{n^2} \right] \right\|_F / \|P^2\|_F \tag{38}$$

The backward error of a solvent $P$, is defined by Higham and Kim (2001) via

$$\eta(P) = \min \left\{ \epsilon : M(P) + \Delta M(P) = 0, \left\| \begin{bmatrix} \Delta A / \|A\|_F & \Delta B / \|B\|_F & \Delta C / \|C\|_F \end{bmatrix} \right\|_F \leq \epsilon \right\} \tag{39}$$

and they show that the backward error is bounded by

$$\underline{\eta(P)} \leq \eta(P) \leq \overline{\eta(P)} \tag{40}$$

where the lower bound on the backward error is

$$\underline{\eta(P)} = \frac{\left\|AP^2 + BP + C\right\|_F}{\left(\left\|P^2\right\|_F^2 / \|A\|_F^2 + \|P\|_F^2 / \|B\|_F^2 + n / \|C\|_F^2\right)^{1/2}} \tag{41}$$

and the upper bound is

$$\overline{\eta(P)} = \frac{\left\|-vec\left(AP^2 + BP + C\right)\right\|_2}{\sigma_{min}\left(\left[\|A\|_F \left(P^2\right)' \otimes I_n \quad \|B\|_F P' \otimes I_n \quad \|C\|_F I_{n^2}\right]\right)} \tag{42}$$

where $vec$ is the columnwise vectorization operator and $\sigma_{min}$ is the smallest singular value. Higham and Kim (2001) point out that their backward error analysis demonstrates that a small relative residual (the absolute residual being $AP^2 + BP + C$ for an approximate $P$ returned by a numerical algorithm) does not necessarily imply a small backward error for the matrix quadratic problem. The importance of this can be seen with the "useful rule of thumb" (Higham, 2002, p. 9)

$$\text{forward error} \lesssim \text{condition number} \times \text{backward error} \tag{43}$$

illustrating that the error in the approximate solution (forward error), necessitates consideration of both the condition number an the backward error. Nonetheless, both bounds on the backward error of solvents are useful when, as is generally the case, an analytic or symbolic solution is not available. The lower bound, $\underline{\eta(P)}$, is particularly appealing to the DSGE literature, which often–e.g., in posterior sampling–requires solving the matrix quadratic equation a multitude of times, as its computational burden is minimal with its calculation via matrix products and norms.

While the backward stability in the calculation of the eigenvalues does not fully characterize the numerical stability of the solution of linear DSGE models, it highlights the potential for QZ or generalized Schur-based algorithms to underperform the alternative algorithms, particularly when the separation between the stable and unstable eigenvalues is small. I will turn to demonstrating this in the next section using the specific model of section 2.

## 5. RESULTS

In this section I investigate the accuracy of the different methods, QZ- and non-QZ-based, from section 3 in solving the model of section 2. The model was chosen to be as

simple as possible, in order to enable the symbolic solution of the underlying matrix quadratic problem; see Higham and Kim (2000) who argue that Matlab can successfully solve two-dimensional matrix quadratic problems reliably. I provide numerical results for two calibrations, see table 1, labeled standard and extreme. The standard calibration follows the RBC literature (see, e.g., King and Rebelo, 1999) with the degree of habit formation, $h$ and curvature in the utility function, $\sigma$, elevated to match an equity premium of 7.8 in annual percentage points following (Mehra, 2003) for the post-war US and $\omega$, the standard deviation of the technology shock, adjusted to deliver a standard deviation of consumption growth, $std(\log c_t)$, of 0.566 in quarterly percent, in line again with the post-war US experience. The extreme calibration is chosen to bring the eigenvalue separation between the stable and unstable pencils closer together, while maintaining the match of the symbolic solution to the equity premium and consumption growth volatility.

Besides assessing whether the different solution methods are able to recover the exact solutions for the two calibration targets, I examine the underlying causes of a degeneration in accuracy following the results of the previous section. Namely the largest absolute deviation in the matrices for the linear solution or policy function (11), $P$ and $Q$, and the largest absolute difference in the finite eigenvalues of the quadratic eigenvalue problem (15) relative to the symbolic solution, and the separation between the calculated stable and unstable eigenvalues along with the conditioning number and bounds on the backward error of the solvent produced by the various methods. Additionally, I provide plots of the pseudospectra of the matrix quadratic (29) and of the QZ companion linearization (32). The results that are referred to as "symbolic" are solved symbolically and evaluated using Mathlab's VPA (variable precision arithmetic) with 100 digits of accuracy.

| | $h$ | $\beta$ | $\delta$ | $\alpha$ | $\sigma$ | $\rho$ | $\omega$ |
|---|---|---|---|---|---|---|---|
| Standard | 0.966 | 0.99 | 0.025 | 0.36 | 98.1 | 0.95 | 0.134 |
| Extreme | 1-3.907E-05 | 1-1.750E-10 | 0.6715 | 1-5.751E-05 | 9.151 | 1-5.184E-04 | 3.068E-03 |

TABLE 1. Calibrations

Table 2 contains the results for the standard calibration. The first line contains the equity premium predicted by the different methods and all of the methods successfully predict an equity premium of 7.8 annual percentage points, likewise the volatility of consumption growth, the third line, is identical across methods. Upon closer examination, the second line, the symbolic equity premium and that predicted by the varying methods

differ to varying degrees. The most accurate methods being those of Binder and Pesaran (1997) and the cyclic reduction method of Dynare, with all QZ-based methods apart from Dynare displaying degrees of accuracy several orders of magnitude lower. As laid out in Villemot (2011), Dynare reduces the problem solved with the QZ algorithm by, among others, eliminating zero column variables in the $A$ and $C$ matrices of the linear system (10); this is in line with one of the suggestions by Hammarling, Munro, and Tisseur (2013) to improve the accuracy of the quadratic eigenvalue problem. This is reflected in the fifth line of the table, where the largest error in the finite eigenvalues calculated by Dynare are in line with the non-QZ-based methods, those of the remaining QZ-based methods are several orders of magnitude larger, and that of Binder and Pesaran (1997) demonstrating the smallest error. The errors in the resulting matrices for the linear solution or policy function (11), $P$ and $Q$ are roughly of the same order of magnitude as the eigenvalue errors. Despite the differences in the accuracy of calculating the eigenvalues, all of the methods yield the same eigenvalue separation between the stable and unstable pencils and conditioning numbers of the solvent $P$. Based on this standard calibration, the differences in the solutions generated by the different methods are of no economic consequence. Yet as indicated by the last two lines, the backward errors in the calculated solvents $P$ differ in a numerically consequential way. Again with the exception of Dynare, the backward errors of all QZ-based methods exceed the double precision unit roundoff of $2^{-52} = 2.2204E - 16$, whereas all other methods deliver a backward error less than the roundoff.

Figure 1 plots the pseudospectra for the extreme standard of the matrix quadratic (29) – in blue – and of the QZ algorithm (32) – in red – against the symbolic eigenvalues – in black – for two different sizes of perturbations. In the left panel, the pseudospectra are not visible, as they overlap with the symbolic results for perturbations of this size. For slightly larger perturbations (right panel), the pseudospectrum of the QZ algorithms encompasses the unit circle while that of the matrix quadratic remains invisible at this scale. This, following Tisseur and Higham (2001), indicates that the backward error in calculating the eigenvalues is not only larger than under the QZ algorithm than with the matrix quadratic, consistent with Tisseur (2000) and with the bounds for the backward errors of the solvents $P$ in table 2, but also that the stable and unstable eigenvalues are potentially indistinguishable numerically.

| Measure | Data | Symbolic | QZ-Based Methods | | | | Alternatives | | |
| --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |
| | | | Klein (2000) | Sims (2001) | Uhlig (1999) | Dynare | Anderson (2010) | BP (1997) | Dynare |
| $E[rp]$ | 7.8 | 7.8 | 7.8 | 7.8 | 7.8 | 7.8 | 7.8 | 7.8 | 7.8 |
| $\Delta E[rp]$ | | | 2.31E-08 | -1.43E-09 | 2.27E-08 | -3.08E-12 | 1.54E-11 | 1.71E-12 | 1.71E-12 |
| $std(\log c_t)$ | 0.566 | 0.566 | 0.566 | 0.566 | 0.566 | 0.566 | 0.566 | 0.566 | 0.566 |
| $\max(|\Delta P|)$ | | | 3.93E-11 | 2.36E-12 | 3.88E-11 | 7.17E-14 | 1.43E-14 | 1.51E-15 | 1.83E-14 |
| $\max(|\Delta Q|)$ | | | 3.21E-11 | 1.99E-12 | 3.14E-11 | 4.25E-15 | 2.13E-14 | 2.40E-15 | 2.39E-15 |
| $\max(|\Delta eig|)$ | | | 4.12E-11 | 2.47E-12 | 4.13E-11 | 1.32E-14 | 2.86E-14 | 3.09E-15 | 1.32E-14 |
| Eig. Sep. | | 0.0127 | 0.0127 | 0.0127 | 0.0127 | 0.0127 | 0.0127 | 0.0127 | 0.0127 |
| $\Psi(P)$ | | 7.27E+03 | 7.27E+03 | 7.27E+03 | 7.27E+03 | 7.27E+03 | 7.27E+03 | 7.27E+03 | 7.27E+03 |
| $\eta(P)$ | | 4.69E-17* | 1.74E-13 | 1.04E-14 | 1.72E-13 | 4.75E-17* | 9.58E-17* | 4.69E-17* | 2.04E-18* |
| $\overline{\eta(P)}$ | | 6.93E-17* | 2.57E-13 | 1.54E-14 | 2.54E-13 | 7.01E-17* | 1.42E-16* | 6.93E-17* | 3.02E-18* |

TABLE 2. Results: Standard Calibration

- For Dynare, refer to Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot (2011). Dynare under QZ-Based Methods is documented in Villemot (2011) and under Alternatives is the cyclic reduction method. BP (1997) refers to Binder and Pesaran (1997).

- $E[rp]$ is expressed in annual %, $std(\log c_t)$ in quarterly %, and * indicates a backward error less than machine precision, $2^{-52} = 2.2204E - 16$.

(A) $1e-6$                                      (B) $2e-5$
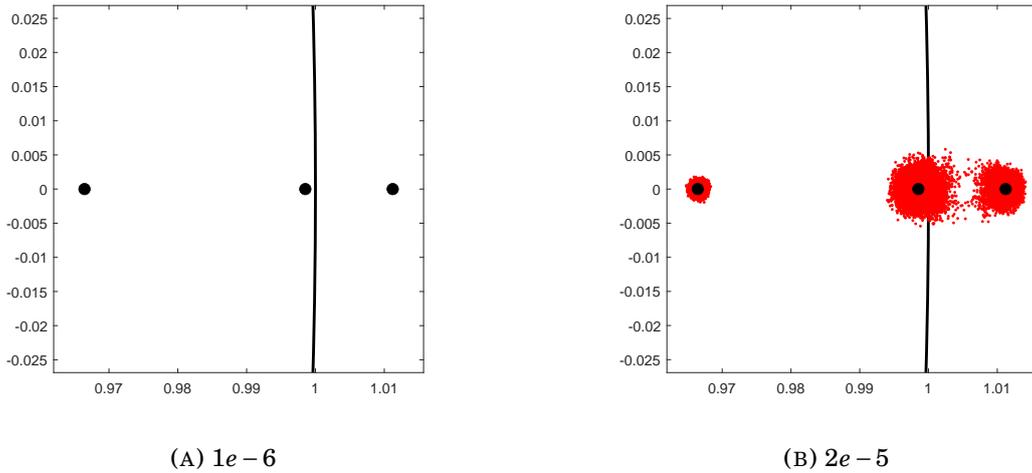
FIGURE 1. Pseudospectrum: Standard Calibration x-axis: real component, y-axis:
imaginary component, large black dots: eigenvalues, black curve: unit circle, small red
dots: pseudospectrum QZ companion linearization (32), small blue dots: pseudospectrum
matrix quadratic (29)

Table 3 contains the results for the extreme calibration and the resulting predictions
for the two calibration targets now differ significantly across methods. While the non-QZ-
based methods continue to maintain a significant match with the calibration targets, lines
1 and 3, the QZ-based methods including Dynare now mispredicts the equity premium
by at least 75 annual basis points and as much as 3 annual percentage points, errors of
genuine economic significance. The second line, containing the differences of the equity
premium predicted by the different methods and the symbolic solution, now show the
algorithm of Anderson (2010) being more accurate than the method of Binder and Pesaran
(1997) and the cyclic reduction method of Dynare being several orders of magnitude
less accurate than either of the two non QZ-based alternatives.. The differences in the
accuracies of the predicted calibration targets ultimately stem from differences in the
accuracies in matrices for the linear solution or policy function (11), $P$ and $Q$, and the
eigenvalues, as can be seen in the third through sixth lines of the table. The eigenvalue
separation has decreased by several orders of magnitude and the conditioning number has
increased by two, again despite differences in others measures, all the algorithms roughly
agree on these two quantities. This is consistent with Kågström and Poromaa's (1996) and
Demmel and Kågström's (1987) assessment that while a generalized eigenvalue problem
is potentially ill-conditioned with respect to eigenvalue groupings if small perturbations
cause eigenvalues of these groupings to coalesce, the eigenvalue separation is in itself
insufficient to determine this and the entire eigenspace must be examined. With respect

(A) $5e - 11$

(B) $2e - 10$
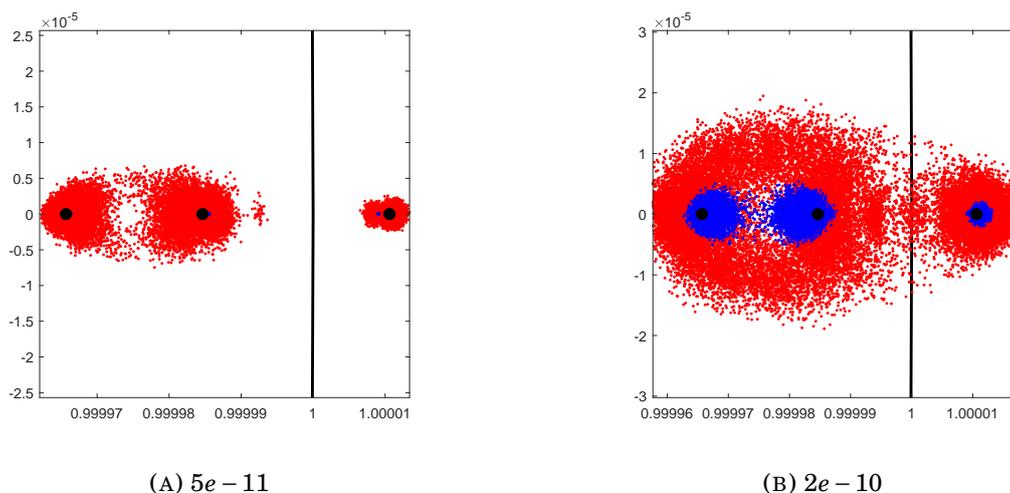
FIGURE 2. Pseudospectrum: Extreme Calibration x-axis: real component, y-axis: imaginary component, large black dots: eigenvalues, black curve: unit circle, small red dots: pseudospectrum QZ companion linearization (32), small blue dots: pseudospectrum matrix quadratic (29)

to the bounds on the backward error, the last two lines of the table now that all QZ-based algorithms and Dynare's cyclic reduction algorithm display backward errors that exceed machine precision. Importantly, not a single one of these algorithms displayed a warning that their solutions might be inaccurate.

Figure 2 plots the pseudospectra for the extreme calibration of the matrix quadratic (29) – in blue – and of the QZ algorithm (32) – in red – against the symbolic eigenvalues – in black – for two different sizes of perturbations. In contrast to the results for the standard calibration in figure 2, the finite eigenvalues are all much closer to the unit circle (see the scale on the x-axis) and dispersion away from the exact eigenvalues is visible with perturbations several orders of magnitude smaller. Again, the pseudospectrum of the QZ algorithm bleeds across the unit circle for smaller perturbations than does the matrix quadratic (right panel).

Table 4 contains a summary of results from additional alternate calibrations (see the appendix, Table 5), in all calibrations, the parameters are chosen to match the annual equity premium of 7.8 and the quarterly standard deviation of consumption growth of 0.566%. Calibrations I and II are alternative "standard" calibrations, holding all parameters apart from $h$, $\sigma$ and $\omega$ constant. Calibration I has a higher curvature in the utility function, $\sigma$, and a lower degree of habit formation, $h$, and calibration II vice versa than in the standard calibration above. As in the standard calibration, the eigenvalue separation is on the order of $1E - 02$ and all methods successfully recover the equity

| Measure | Data | Symbolic | QZ-Based Methods | | | | Alternatives | | |
|---|---|---|---|---|---|---|---|---|---|
| | | | Klein (2000) | Sims (2001) | Uhlig (1999) | Dynare | Anderson (2010) | BP (1997) | Dynare |
| $E[rp]$ | 7.8 | 7.8 | 7.05 | 4.75 | 7.05 | 6.47 | 7.8 | 7.8 | 7.8 |
| $\Delta E[rp]$ | | | 0.754 | 3.05 | 0.754 | 1.33 | -8.31E-07 | -3.06E-06 | 0.000143 |
| $std(\log c_t)$ | 0.566 | 0.566 | 0.529 | 0.41 | 0.529 | 0.491 | 0.567 | 0.567 | 0.567 |
| $\max(|\Delta P|)$ | | | 1.01E-06 | 4.07E-06 | 1.01E-06 | 1.74E-06 | 1.13E-12 | 4.01E-12 | 1.87E-10 |
| $\max(|\Delta Q|)$ | | | 0.00127 | 0.00515 | 0.00127 | 0.00224 | 1.40E-09 | 5.16E-09 | 2.41E-07 |
| $\max(|\Delta eig|)$ | | | 1.84E-06 | 6.51E-06 | 1.84E-06 | 5.73E-06 | 1.55E-12 | 1.57E-11 | 5.73E-06 |
| Eig. Sep. | | 2.82E-05 | 2.60E-05 | 2.09E-05 | 2.60E-05 | 2.55E-05 | 2.82E-05 | 2.82E-05 | 2.55E-05 |
| $\Psi(P)$ | | 5.36E+05 | 5.32E+05 | 5.07E+05 | 5.32E+05 | 5.07E+05 | 5.36E+05 | 5.36E+05 | 5.36E+05 |
| $\eta(P)$ | | 3.59E-17* | 6.73E-12 | 3.53E-11 | 6.73E-12 | 4.70E-11 | 3.68E-17* | 1.43E-16* | 4.59E-15 |
| $\overline{\eta(P)}$ | | 5.07E-17* | 9.51E-12 | 4.99E-11 | 9.51E-12 | 6.65E-11 | 5.20E-17* | 2.03E-16* | 6.49E-15 |

TABLE 3. Results: Extreme Calibration

- For Dynare, refer to Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot (2011). Dynare under QZ-Based Methods is documented in Villemot (2011) and under Alternatives is the cyclic reduction method. BP (1997) refers to Binder and Pesaran (1997).

- $E[rp]$ is expressed in annual %, $std(\log c_t)$ in quarterly %, and $*$ indicates a backward error less than machine precision, $2^{-52} = 2.2204E - 16$.

| Calibration | Symbolic | QZ-Based Methods | | | | Alternatives | | |
| --- | --- | --- | --- | --- | --- | --- | --- | --- |
| | | Klein (2000) | Sims (2001) | Uhlig (1999) | Dynare | Anderson (2010) | BP (1997) | Dynare |
| I | 0.011 | -9.41E-14 | 4.85E-11 | -3.98E-12 | 6.49E-10 | -5.03E-13 | 5.42E-13 | 1.24E-12 |
| | 9.69E+03 | 5.23E-17* | 2.28E-16 | 3.99E-16 | 8.04E-14 | 2.18E-17* | 3.12E-18* | 5.00E-17* |
| II | 0.0105 | 3.09E-10 | 1.49E-09 | 2.97E-10 | -5.61E-10 | -1.72E-12 | 1.02E-11 | 1.52E-11 |
| | 1.50E+05 | 1.56E-15 | 1.04E-16* | 1.76E-15 | 4.00E-16 | 5.44E-18* | 1.94E-19* | 4.97E-17* |
| III | 0.000117 | -0.505 | -0.373 | -0.505 | 1.3 | 3.02E-07 | 3.45E-07 | 9.96E-05 |
| | 1.01E+06 | 4.83E-11 | 3.73E-11 | 4.83E-11 | 1.04E-10 | 9.48E-17* | 4.65E-17* | 5.09E-14 |
| IV | 1.97E-05 | -4.9 | -0.176 | -4.9 | 4.61E-05 | -1.69E-05 | -4.38E-06 | -4.59E-06 |
| | 3.24E+05 | 1.85E-11 | 2.74E-11 | 1.85E-11 | 2.46E-16 | 8.98E-17* | 8.11E-17* | 1.82E-17* |
| V | 3.20E-05 | 0.0631 | -3.91 | 0.0568 | -9.19E-06 | -5.46E-06 | -4.50E-06 | -2.23E-05 |
| | 4.50E+05 | 1.23E-11 | 5.76E-11 | 1.23E-11 | 6.93E-17* | 5.55E-17* | 5.10E-17* | 7.65E-16 |
| VI | 2.90E-05 | 0.812 | -1.84 | 0.813 | 2.28E-05 | 4.24E-06 | -5.78E-06 | 2.92E-05 |
| | 3.89E+05 | 1.52E-11 | 2.37E-11 | 1.52E-11 | 4.29E-16 | 4.92E-17* | 6.02E-17* | 9.64E-16 |

TABLE 4. Results for additional calibrations I-VI (see the appendix), satisfying $E[rp] = 7.8$ and $std(\log c_t) = 0.566\%$ symbolically. For the column "Symbolic", the first row is $Eig.Sep.$, the second $\Psi(P)$. For all other columns, the first row is $E[\Delta rp]$, the second $\overline{\eta(P)}$. For Dynare, refer to Adjemian, Bastani, Juillard, Mihoubi, Perendia, Ratto, and Villemot (2011). Dynare under QZ-Based Methods is documented in Villemot (2011) and under Alternatives is the cyclic reduction method. BP (1997) refers to Binder and Pesaran (1997).

* indicates a backward error less than machine precision, $2^{-52} = 2.2204E - 16$.

premium. Calibration III is similar to the extreme calibration above, but with a slightly reduced degree of habit formation, $h$, and discount factor, $\beta$, compensated by an increased curvature in the utility function, $\sigma$. With an eigenvalue separation on the order of $1E-04$, the QZ methods demonstrate significant deviations in their predicted equity premia as above, though now some methods over and some methods under predict the premium. Both the alternative methods successfully match the premium; under this calibration, the accuracy of the Binder and Pesaran (1997) method exceeds that of Anderson (2010) (not shown in the table, full results are available online). Calibrations IV-VI provide further examples of potentially arbitrary results from QZ methods. With the eigenvalue separation on the order of $1E-05$, some methods do very well for some calibrations yet worse for others, with all the calibrated parameters arguably very similar (see the appendix, Table 5). Interestingly, Dynare in all of these latter three calibrations arguably predicts the equity premium successfully. In terms of the backward errors, only Anderson (2010) and Binder and Pesaran (1997) produce solvents, $P$, with errors always less than the unit roundoff. Although Dynare's QZ method only produces errors below this threshold for calibration V, its errors are on the same order of magnitude for calibrations II, IV, and VI. Likewise Dynare's cyclic reduction method fairs better than the QZ algorithms, albeit not at quite the level of the other two alternatives to QZ. Again, none of the algorithms produced any warning as to the potential inaccuracy of their solutions.

## 6. Conclusion

This paper has provided a concrete example, calibrated to macroeconomic and financial data, of economically significant numerical errors from standard QZ-based methods. In this example, the errors become significant when the DSGE model has insufficient separation between the backward looking, or stable, and the forward looking, or unstable, components of the solution. None of the QZ-based methods from the literature examined here gave the user any indication that the numerical solution it provided might imply an equity premium off by several percentage points. This example was chosen to be as small as possible, two endogenous variables, to enable a benchmark symbolic solution, but backward error analyses successfully warn of the inaccuracies without appealing to this symbolic solution. Thus, this paper serves as a cautionary tale and provides a first step towards numerical accuracy metrics for the ubiquitous linear DSGE model.

REFERENCES

ADJEMIAN, S., H. BASTANI, M. JUILLARD, F. MIHOUBI, G. PERENDIA, M. RATTO, AND S. VILLEMOT (2011): "Dynare: Reference Manual, Version 4," Dynare Working Papers 1, CEPREMAP.

ANDERSON, G. S. (2008): "Solving Linear Rational Expectations Models: A Horse Race," *Computational Economics*, 31(2), 95–113.

——— (2010): "A Reliable and Computationally Efficient Algorithm for Imposing the Saddle Point Property in Dynamic Models," *Journal of Economic Dynamics and Control*, 34(3), 472–489.

ANDERSON, G. S., A. LEVIN, AND E. SWANSON (2006): "Higher-Order Pertubation Solutions to Dynamic Discrete-Time Rational Expectations Models," Discussion Paper 2006-01, Federal Reserve Bank of San Francisco Working Paper Series.

ANDERSON, G. S., AND G. MOORE (1985): "A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models," *Economics Letters*, 17(3), 247–252.

ANGUAS, L. M., M. I. BUENO, AND F. M. DOPICO (2019): "A comparison of eigenvalue condition numbers for matrix polynomials," *Linear Algebra and its Applications*, 564, 170 − 200.

BINDER, M., AND M. H. PESARAN (1997): "Multivariate Linear Rational Expectations Models: Characterization of the Nature of the Solutions and Their Fully Recursive Computation," *Econometric Theory*, 13(6), 877–88.

BINI, D. A., G. LATOUCHE, AND B. MEINI (2002): "Solving matrix polynomial equations arising in queueing problems," *Linear Algebra and its Applications*, 340(1), 225–244.

BLANCHARD, O. J. (1979): "Backward and Forward Solutions for Economies with Rational Expectations," *The American Economic Review*, 69(2), 114–118.

BLANCHARD, O. J., AND C. M. KAHN (1980): "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 48(5), 1305–1311.

CAMPBELL, J. Y. (2003): "Consumption-based asset pricing," in *Financial Markets and Asset Pricing*, vol. 1 of *Handbook of the Economics of Finance*, chap. 13, pp. 803–887. Elsevier.

CAMPBELL, J. Y., AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

CAMPBELL, J. Y., AND R. J. SHILLER (1988): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *The Review of Financial Studies*, 1(3), 195–228.

COCHRANE, J. H. (2008): "Financial Markets and the Real Economy," in *Handbook of the Equity Risk Premium*, ed. by R. Mehra, Handbooks in Finance, pp. 237 – 325. Elsevier, San Diego.

CONSTANTINIDES, G. M. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98(3), 519–543.

CROCE, M. M. (2014): "Long-run productivity risk: A new hope for production-based asset pricing?," *Journal of Monetary Economics*, 66(C), 13–31.

DEMMEL, J. W., AND B. KÅGSTRÖM (1987): "Computing Stable Eigendecompositions of Matrix Pencils," *Linear Algebra and its Applications*, (88/89), 139–186.

DENNIS, JR., J. E., J. F. TRAUB, AND R. P. WEBER (1976): "The Algebraic Theory of Matrix Polynomials," *SIAM Journal on Numerical Analysis*, 13(6), 831–845.

FERNÁNDEZ-VILLAVERDE, J., J. RUBIO-RAMÍREZ, AND F. SCHORFHEIDE (2016): "Solution and Estimation Methods for DSGE Models," in *Handbook of Macroeconomics, Volume 2A*, ed. by J. B. Taylor, and H. Uhlig, Handbooks in Economics, chap. 9, pp. 527–724. North-Holland/Elsevier.

GANTMACHER, F. R. (1959): *The Theory of Matrices*, vol. I&II. Chelsea Publishing Company, New York, NY.

GOLUB, G. H., AND C. F. VAN LOAN (2013): *Matrix Computations*. The Johns Hopkins University Press, fourth edn.

HAMMARLING, S., C. J. MUNRO, AND F. TISSEUR (2013): "An Algorithm for the Complete Solution of Quadratic Eigenvalue Problems," *ACM Transactions On Mathematical Software*, 39(3), 18:1–18:19.

HANSEN, L. P., AND K. J. SINGLETON (1983): "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91(2), 249–265.

HEILBERGER, C., T. KLARL, AND A. MAUSSNER (2015): "On the uniqueness of solutions to rational expectations models," *Economics Letter*, 128, 14–16.

HIGHAM, N. J. (2002): *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, second edn.

HIGHAM, N. J., AND H.-M. KIM (2000): "Numerical Analysis of a Quadratic Matrix Equation," *IMA Journal of Numerical Analysis*, 20, 499–519.

———— (2001): "Solving a Quadratic Matrix Equation by Newton's Method with Exact Line Searches," *SIAM Journal on Matrix Analysis and Applications*, 23(2), 499–519.

HIGHAM, N. J., D. S. MACKEY, AND F. TISSEUR (2006): "The Conditioning of Linearizations of Matrix Polynomials," *SIAM Journal on Matrix Analysis and Applications*, 28(4), 1005–1028.

HIGHAM, N. J., D. S. MACKEY, F. TISSEUR, AND S. D. GARVEY (2008): "Scaling, sensitivity and stability in the numerical solution of quadratic eigenvalue problems," *International Journal for Numerical Methods in Engineering*, 73(3), 344–360.

HIGHAM, N. J., AND F. TISSEUR (2002): "More on pseudospectra for polynomial eigenvalue problems and applications in control theory," *Linear Algebra and its Applications*, 351-352, 435 – 453, Fourth Special Issue on Linear Systems and Control.

JERMANN, U. J. (1998): "Asset Pricing in Production Economies," *Journal of Monetary Economics*, 41(2), 257–275.

JIN, H.-H., AND K. L. JUDD (2002): "Pertubation Methods for General Dynamic Stochastic Models," Mimeo December, Stanford University.

JUDD, K. L. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, MA.

KÅGSTRÖM, B., AND P. POROMAA (1996): "Computing eigenspaces with specified eigenvalues of a regular matrix pair (A,B) and condition estimation: theory, algorithms and software," *Numerical Algorithms*, (12), 369–407.

KING, R. G., AND S. T. REBELO (1999): "Resuscitating real business cycles," vol. 1 of *Handbook of Macroeconomics*, chap. 14, pp. 927–1007. Elsevier.

KLEIN, P. (2000): "Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model," *Journal of Economic Dynamics and Control*, 24(10), 1405–1423.

KYDLAND, F. E., AND E. C. PRESCOTT (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50(6), 1345–1370.

LAN, H., AND A. MEYER-GOHDE (2014): "Solvability of Perturbation Solutions in DSGE Models," *Journal of Economic Dynamics and Control*, 45, 366–388.

LETTAU, M. (2003): "Inspecting the Mechanism: Closed-Form Solutions for Asset Prices in Real Business Cycle Models," *The Economic Journal*, 113(489), 550–575.

LUENBERGER, D. G. (1978): "Time-invariant descriptor systems," *Automatica*, 14(5), 473–480.

MEHRA, R. (2003): "The Equity Premium: Why Is It a Puzzle?," *Financial Analysts Journal*, 59(1), 54–69.

MEHRA, R., AND E. C. PRESCOTT (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15(2), 145–161.

MENGI, E., AND M. L. OVERTON (2005): "Algorithms for the computation of the pseudospectral radius and the numerical radius of a matrix," *IMA Journal of Numerical Analysis*, 25(4), 648–669.

MICHIELS, W., K. GREEN, T. WAGENKNECHT, AND S.-I. NICULESCU (2006): "Pseudospectra and stability radii for analytic matrix functions with application to time-delay systems," *Linear Algebra and its Applications*, 418(1), 315 – 335.

MOLER, C. B., AND G. W. STEWART (1973): "An Algorithm for Generalized Matrix Eigenvalue Problems," *SIAM Journal on Numerical Analysis*, 10(2), 241–256.

SIMS, C. A. (2001): "Solving Linear Rational Expectations Models," *Computational Economics*, 20(1-2), 1–20.

STEWART, G. W. (1972): "On the Sensitivity of the Eigenvalue Problem $Ax = \lambda Bx$," *SIAM Journal on Numerical Analysis*, 9(4), 669–686.

TALLARINI, JR., T. D. (2000): "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 45(3), 507–532.

TISSEUR, F. (2000): "Backward error and condition of polynomial eigenvalue problems," *Linear Algebra and its Applications*, 309(1), 339–361.

TISSEUR, F., AND N. J. HIGHAM (2001): "Structured Pseudospectra for Polynomial Eigenvalue Problems, with Applications," *SIAM Journal on Matrix Analysis and Applications*, 23(1), 187–208.

TISSEUR, F., AND K. MEERBERGEN (2001): "The Quadratic Eigenvalue Problem," *SIAM Review*, 43(2), 235–286.

UHLIG, H. (1999): "A Toolkit for Analysing Nonlinear Dynamic Stochastic Models Easily," in *Computational Methods for the Study of Dynamic Economies*, ed. by R. Marimon, and A. Scott, chap. 3, pp. 30–61. Oxford University Press.

VILLEMOT, S. (2011): "Solving Rational Expectations Models at First Order: What Dynare Does," Dynare Working Papers 2, CEPREMAP.

WILKINSON, J. (1979): "Kronecker's canonical form and the QZ algorithm," *Linear Algebra and its Applications*, 28, 285 − 303.

## APPENDIX

6.1. **Multivariate pivot derivation of the linear solution using the generalized Schur decomposition.** While this derivation contains nothing substantially new compared with, say Klein (2000), its formulation commensurate with (10) enables a straightforward application of Blanchard's (1979) forward method, making the derivations potentially more transparent and accessible than existing expositions in the literature.

Rearranging the model (10) into the companion linearization yields

$$F \begin{bmatrix} y_t \\ E_t[y_{t+1}] \end{bmatrix} = G \begin{bmatrix} y_{t-1} \\ y_t \end{bmatrix} + \begin{bmatrix} 0_{n_y \times n_\varepsilon} \\ D \end{bmatrix} \varepsilon_t, \quad F \equiv \begin{bmatrix} I_{n_y} & 0_{n_y \times n_y} \\ 0_{n_y \times n_y} & A \end{bmatrix}, G \equiv \begin{bmatrix} 0_{n_y \times n_y} & I_{n_y} \\ -C & -B \end{bmatrix} \tag{A1}$$

where $I_{n_y}$ is an $n_y \times n_y$ identity matrix and $0_{n_y \times n_y}$ is an $n_y \times n_y$ zero matrix.

The generalized Schur decomposition (unitary $Q$ and $Z$ and upper triangular $S$ and $T$ with $Q^* F Z = S$ and $Q^* G Z = T$) of the matrix pencil $P_{FG}(z) = Fz - G$, can be ordered arbitrarily to form

$$\begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} \begin{bmatrix} E_t[w^s_{t+1}] \\ E_t[w^u_{t+1}] \end{bmatrix} = \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix} \begin{bmatrix} w^s_t \\ w^u_t \end{bmatrix} + Q^* \begin{bmatrix} 0_{n_y \times n_\varepsilon} \\ D \end{bmatrix} \varepsilon_t \tag{A2}$$

with the definition $Z \begin{bmatrix} w^{s\prime}_t & w^{u\prime}_t \end{bmatrix}' = \begin{bmatrix} y'_{t-1} & y'_t \end{bmatrix}'$. With any generalized Schur decomposition of $P_{DE}(z)$, the spectrum or set of eigenvalues of the pencil $P_{DE}(z)$ is determined by the diagonal entries of $S$ and $T$

$$\rho(P_{DE}) = \{t_{ii}/s_{ii}, \text{ if } s_{ii} \neq 0; \infty, \text{ if } s_{ii} = 0; \emptyset, \text{ if } s_{ii} = t_{ii} = 0; i = 1, \ldots, 2n_y\} \tag{A3}$$

where $s_{ii}$ and $t_{ii}$ denote the $i$'th row and $i$'th column of $S$ and $T$ respectively. Ordering the decomposition so that the unstable eigenvalues are in the lower right blocks of $S$ and $T$ (hence $S_{22}$ and $T_{22}$), this lower block can be solved forward following Blanchard (1979) to yield

$$w^u_t = \lim_{j \to \infty} \left(T_{22}^{-1} S_{22}\right)^j E_t\left[w^u_{t+j}\right] - T_{22}^{-1} \underbrace{\begin{bmatrix} \{Q^*\}_{21} & \{Q^*\}_{22} \end{bmatrix}}_{\equiv \{Q^{-1}\}_{2\bullet}} \underbrace{\begin{bmatrix} 0'_{n_y \times n_\varepsilon} & D' \end{bmatrix}'}_{\equiv \hat{D}} \varepsilon_t = -T_{22}^{-1} \{Q^*\}_{2\bullet} \hat{D} \varepsilon_t \tag{A4}$$

where the invertibility of $T_{22}$ and convergence of $\lim_{j \to \infty} \left(T_{22}^{-1} S_{22}\right)^j$ follow directly from the ordering above. Using the definition $Z \begin{bmatrix} w^{s\prime}_t & w^{u\prime}_t \end{bmatrix}' = \begin{bmatrix} y'_{t-1} & y'_t \end{bmatrix}'$ from above delivers

$$w^u_t = \begin{bmatrix} Z^*_{21} & Z^*_{22} \end{bmatrix} \begin{bmatrix} y'_{t-1} & y'_t \end{bmatrix}' = -T_{22}^{-1} \{Q^*\}_{2\bullet} \hat{D} \varepsilon_t \tag{A5}$$

where $^*$ indicates the complex conjugation of $Z$ that delivers its inverse by virtue of it being a unitary matrix. If the necessary and sufficient assumptions for a unique stable solution for $y_t$ of (10) from the main text hold, the unique stable solution for $y_t$ is given by

$$y_t = Z_{21} Z_{11}^{-1} y_{t-1} - \left(Z_{22} - Z_{21} Z_{11}^{-1} Z_{12}\right) T_{22}^{-1} \{Q^*\}_{2\bullet} \hat{D} \varepsilon_t \tag{A6}$$

$$= Q_{11} S_{11}^{-1} T_{11} Q_{11}^{-1} y_{t-1} - \left(Z_{22} - Z_{21} Z_{11}^{-1} Z_{12}\right) T_{22}^{-1} \{Q^*\}_{2\bullet} \hat{D} \varepsilon_t \tag{A7}$$

where $Z_{22}^{*}{}^{-1} = Z_{22} - Z_{21}Z_{11}^{-1}Z_{12}$ and $Z_{22}^{*}{}^{-1}Z_{21}^{*} = -Z_{21}Z_{11}^{-1}$ follow from the properties of unitary matrices and $Z_{21}Z_{11}^{-1} = Q_{11}S_{11}^{-1}T_{11}Q_{11}^{-1}$ from the first block rows of $F$ and $G$ in (16) and upper triangularity of $S$ and $T$. From $Q_{11}S_{11}^{-1}T_{11}Q_{11}^{-1}$, it follows that the recursion in $y_t$ is stable from the ordering of the eigenvalues above, i.e. the eigenvalues of the upper left block of the generalized Schur decomposition, $det(S_{11}\lambda - T_{11}) = 0$, are inside the unit circle.

|      | $h$          | $\beta$       | $\delta$  | $\alpha$      | $\sigma$      | $\rho$       | $\omega$   |
|------|--------------|---------------|-----------|---------------|---------------|--------------|------------|
| I    | 0.8617       | 0.99          | 0.025     | 0.36          | 324.3         | 0.95         | 8.355E-02  |
| II   | 1-9.857E-05  | 0.99          | 0.025     | 0.36          | 6.109         | 0.95         | 6.175E-02  |
| III  | 1-1.008E-04  | 1-8.991E-06   | 0.6402    | 1-5.680E-04   | 51.53         | 1-6.066E-05  | 7.742E-04  |
| IV   | 1-6.829E-06  | 1-5.863E-08   | 0.6562    | 1-2.652E-05   | 1+2.591E-08   | 1-3.437E-03  | 1.594E-02  |
| V    | 1-4.294E-06  | 1-1.012E-12   | 0.4727    | 1-9.990E-05   | 1+7.590E-08   | 1-9.628E-04  | 7.898E-03  |
| VI   | 1-5.070E-06  | 1-4.259E-08   | 0.6539    | 1-5.715E-05   | 1+4.755E-05   | 1-1.221E-03  | 7.102E-03  |

TABLE 5. Additional Calibrations I-VI