Author's personal copy

Review of Economic Dynamics 15 (2012) 188–206

Contents lists available at SciVerse ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

Computing DSGE models with recursive preferences and stochastic volatility ☆

Dario Caldara [a], Jesús Fernández-Villaverde [b,c,d,e,*], Juan F. Rubio-Ramírez [f,g,e], Wen Yao [b]

[a] Federal Reserve Board, Washington, DC 20551, United States
[b] University of Pennsylvania, 160 McNeil, 3718 Locust Walk, Philadelphia, PA 19104, United States
[c] NBER, United States
[d] CEPR, United Kingdom
[e] FEDEA, Spain
[f] Duke University, 213 Social Sciences Building, Box 90097, Durham, NC 27708-0097, United States
[g] Federal Reserve Bank of Atlanta, United States

ARTICLE INFO

Article history:
Received 20 September 2011
Available online 17 October 2011

JEL classification:
C63
C68
E32

Keywords:
DSGE models
Recursive preferences
Perturbation

ABSTRACT

This paper compares different solution methods for computing the equilibrium of dynamic stochastic general equilibrium (DSGE) models with recursive preferences such as those in Epstein and Zin (1989, 1991) and stochastic volatility. Models with these two features have recently become popular, but we know little about the best ways to implement them numerically. To fill this gap, we solve the stochastic neoclassical growth model with recursive preferences and stochastic volatility using four different approaches: second- and third-order perturbation, Chebyshev polynomials, and value function iteration. We document the performance of the methods in terms of computing time, implementation complexity, and accuracy. Our main finding is that perturbations are competitive in terms of accuracy with Chebyshev polynomials and value function iteration while being several orders of magnitude faster to run. Therefore, we conclude that perturbation methods are an attractive approach for computing this class of problems.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

This paper compares different solution methods for computing the equilibrium of dynamic stochastic general equilibrium (DSGE) models with recursive preferences and stochastic volatility (SV). Both features have become very popular in finance and in macroeconomics as modeling devices to account for business cycle fluctuations and asset pricing. Recursive preferences, as those first proposed by Kreps and Porteus (1978) and later generalized by Epstein and Zin (1989, 1991) and Weil (1990), are attractive for two reasons. First, they allow us to separate risk aversion and intertemporal elasticity of substitution (EIS). Second, they offer the intuitive appeal of having preferences for early or later resolution of uncertainty (see the reviews by Backus et al., 2004, 2007, and Hansen et al., 2007, for further details and references). SV generates heteroskedastic aggregate fluctuations, a basic property of many time series such as output (see the review by Fernández-Villaverde and Rubio-Ramírez, 2010), and adds extra flexibility in accounting for asset pricing patterns. In fact, in an influential paper,

☆ We thank Michel Juillard for his help with computational issues and Larry Christiano, Dirk Krueger, Pawel Zabczyk, and participants at several seminars for comments. Beyond the usual disclaimer, we must note that any views expressed herein are those of the authors and not necessarily those of the Board of Governors of the Federal Reserve System or the Federal Reserve Bank of Atlanta. Finally, we also thank the NSF for financial support.

* Corresponding author at: University of Pennsylvania, 160 McNeil, 3718 Locust Walk, Philadelphia, PA 19104, United States.
E-mail addresses: dario.caldara@frb.gov (D. Caldara), jesusfv@econ.upenn.edu (J. Fernández-Villaverde), juan.rubio-ramirez@duke.edu (J.F. Rubio-Ramírez), wenyao@econ.upenn.edu (W. Yao).

1094-2025/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.red.2011.10.001

Author's personal copy

Bansal and Yaron (2004) have argued that the *combination* of recursive preferences and SV is the key for their proposed mechanism, long-run risk, to be successful at explaining asset pricing.

But despite the popularity and importance of these issues, nearly nothing is known about the numerical properties of the different solution methods that solve equilibrium models with recursive preferences and SV. For example, we do not know how well value function iteration (VFI) performs or how good local approximations are compared with global ones. Similarly, if we want to estimate the model, we need to assess which solution method is sufficiently reliable yet quick enough to make the exercise feasible. More important, the most common solution algorithm in the DSGE literature, (log-) linearization, cannot be applied, since it makes us miss the whole point of recursive preferences or SV: the resulting (log-) linear decision rules are certainty equivalent and do not depend on risk aversion or volatility. This paper attempts to fill this gap in the literature, and therefore, it complements previous work by Aruoba et al. (2006), in which a similar exercise is performed with the neoclassical growth model with CRRA utility function and constant volatility.

We solve and simulate the model using four main approaches: perturbation (of second and third order), Chebyshev polynomials, and VFI. By doing so, we span most of the relevant methods in the literature. Our results provide a strong guess of how some other methods not covered here, such as finite elements, would work (rather similar to Chebyshev polynomials but more computationally intensive). We report results for a benchmark calibration of the model and for alternative calibrations that change the variance of the productivity shock, the risk aversion, and the intertemporal elasticity of substitution. In that way, we study the performance of the methods both for cases close to the CRRA utility function with constant volatility and for highly non-linear cases far away from the CRRA benchmark. For each method, we compute decision rules, the value function, the ergodic distribution of the economy, business cycle statistics, the welfare costs of aggregate fluctuations, and asset prices. Also, we evaluate the accuracy of the solution by reporting Euler equation errors.

We highlight four main results from our exercise. First, all methods provide a high degree of accuracy. Thus, researchers who stay within our set of solution algorithms can be confident that their quantitative answers are sound.

Second, perturbations deliver a surprisingly high level of accuracy with considerable speed. Both second- and third-order perturbations perform remarkably well in terms of accuracy for the benchmark calibration, being competitive with VFI or Chebyshev polynomials. For this calibration, a second-order perturbation that runs in a fraction of a second does nearly as well in terms of the average Euler equation error as a VFI that takes ten hours to run. Even in the extreme calibration with high risk aversion and high volatility of productivity shocks, perturbation works at a more than acceptable level. Since, in practice, perturbation methods are the only computationally feasible method to solve the medium-scale DSGE models used for policy analysis that have dozens of state variables (as in Smets and Wouters, 2007), this finding has an outmost applicability. Moreover, since implementing second- and third-order perturbations is feasible with off-the-shelf software like `Dynare`, which requires minimum programming knowledge by the user, our findings may induce many researchers to explore recursive preferences and/or SV in further detail. Two final advantages of perturbation are that, often, the perturbed solution provides insights about the economics of the problem and that it might be an excellent initial guess for VFI or for Chebyshev polynomials.

Third, Chebyshev polynomials provide a terrific level of accuracy with reasonable computational burden. When accuracy is most required and the dimensionality of the state space is not too high, as in our model, they are the obvious choice.

Fourth, we were disappointed by the poor performance of VFI, which, compared with Chebyshev, could not achieve a high accuracy even with a large grid. This suggests that we should relegate VFI to solving those problems where non-differentiabilities complicate the application of the previous methods.

The rest of the paper is organized as follows. In Section 2, we present our test model. Section 3 describes the different solution methods used to approximate the decision rules of the model. Section 4 discusses the calibration of the model. Section 5 reports numerical results and Section 6 concludes. Appendix A provides some additional details.

2. The stochastic neoclassical growth model with recursive preferences and SV

We use the stochastic neoclassical growth model with recursive preferences and SV in the process for technology as our test case. We select this model for three reasons. First, it is the workhorse of modern macroeconomics. Even more complicated New Keynesian models with real and nominal rigidities, such as those in Woodford (2003) or Christiano et al. (2005), are built around the core of the neoclassical growth model. Thus, any lesson learned with it is likely to have a wide applicability. Second, the model is, except for the form of the utility function and the process for SV, the same test case as in Aruoba et al. (2006). This provides us with a set of results to compare to our findings. Three, the introduction of recursive preferences and SV make the model both more non-linear (and hence, a challenge for different solution algorithms) and potentially more relevant for practical use. For example, and as mentioned in the Introduction, Bansal and Yaron (2004) have emphasized the importance of the combination of recursive preferences and time-varying volatility to account for asset prices.

The description of the model is straightforward, and we just go through the details required to fix notation. There is a representative household that has preferences over streams of consumption, c_t, and leisure, $1-l_t$, represented by a recursive function of the form:

$$U_t = \max_{c_t, l_t}\Big[(1-\beta)\big(c_t^{\upsilon}(1-l_t)^{1-\upsilon}\big)^{\frac{1-\gamma}{\theta}} + \beta\big(\mathbb{E}_t U_{t+1}^{1-\gamma}\big)^{\frac{1}{\theta}}\Big]^{\frac{\theta}{1-\gamma}}. \tag{1}$$

Author's personal copy

The parameters in these preferences include β, the discount factor, υ, which controls labor supply, γ, which controls risk aversion, and:

$$\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$$

where ψ is the EIS. The parameter θ is an index of the deviation with respect to the benchmark CRRA utility function (when $\theta = 1$, we are back in that CRRA case where the inverse of the EIS and risk aversion coincide).

The household's budget constraint is given by:

$$c_t + i_t + \frac{b_{t+1}}{R_t^f} = w_t l_t + r_t k_t + b_t$$

where i_t is investment, R_t^f is the risk-free gross interest rate, b_t is the holding of an uncontingent bond that pays 1 unit of consumption good at time $t+1$, w_t is the wage, l_t is labor, r_t is the rental rate of capital, and k_t is capital. Asset markets are complete and we could have also included in the budget constraint the whole set of Arrow securities. Since we have a representative household, this is not necessary because the net supply of any security is zero. Households accumulate capital according to the law of motion $k_{t+1} = (1-\delta)k_t + i_t$ where δ is the depreciation rate.

The final good in the economy is produced by a competitive firm with a Cobb–Douglas technology $y_t = e^{z_t} k_t^{\zeta} l_t^{1-\zeta}$ where z_t is the productivity level that follows:

$$z_t = \lambda z_{t-1} + e^{\sigma_t}\varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1).$$

Stationarity is the natural choice for our exercise. If we had a deterministic trend, we would only need to adjust β in our calibration below (and the results would be nearly identical). If we had a stochastic trend, we would need to rescale the variables by the productivity level and solve the transformed problem. However, in this case, it is well known that the economy fluctuates less than when $\lambda < 1$, and therefore, all solution methods would be closer, limiting our ability to appreciate differences in their performance.

The innovation ε_t is scaled by an SV level σ_t, which evolves as:

$$\sigma_t = (1-\rho)\overline{\sigma} + \rho\sigma_{t-1} + \eta\omega_t, \quad \omega_t \sim \mathcal{N}(0,1)$$

where $\overline{\sigma}$ is the unconditional mean level of σ_t, ρ is the persistence of the processes, and η is the standard deviation of the innovations to σ_t. Our specification is parsimonious and it introduces only two new parameters, ρ and η. At the same time, it captures some important features of the data (see a detailed discussion in Fernández-Villaverde and Rubio-Ramírez, 2010). The combination of an exponent in the productivity process (e^{σ_t}) and a level in the evolution of σ_t generates interesting non-linear dynamics. Another important point is that, with SV, we have two innovations, an innovation to technology, ε_t, and an innovation to the standard deviation of technology, ω_t. Finally, the economy must satisfy the aggregate resource constraint $y_t = c_t + i_t$.

The definition of equilibrium is standard and we skip it in the interest of space. Also, both welfare theorems hold, a fact that we will exploit by jumping back and forth between the solution of the social planner's problem and the competitive equilibrium. However, this is only to simplify our derivations. It is straightforward to adapt the solution methods described below to solve problems that are not Pareto optimal.

Thus, an alternative way to write this economy is to look at the value function representation of the social planner's problem in terms of its three state variables, capital k_t, productivity z_t, and volatility, σ_t:

$$\begin{aligned}
V(k_t, z_t, \sigma_t) = \max_{c_t, l_t} & \left[(1-\beta)\left(c_t^{\upsilon}(1-l_t)^{1-\upsilon}\right)^{\frac{1-\gamma}{\theta}} + \beta\left(\mathbb{E}_t V^{1-\gamma}(k_{t+1}, z_{t+1}, \sigma_{t+1})\right)^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}} \\
\text{s.t.} \quad & c_t + k_{t+1} = e^{z_t} k_t^{\zeta} l_t^{1-\zeta} + (1-\delta)k_t \\
& z_t = \lambda z_{t-1} + e^{\sigma_t}\varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1) \\
& \sigma_t = (1-\rho)\overline{\sigma} + \rho\sigma_{t-1} + \eta\omega_t, \quad \omega_t \sim \mathcal{N}(0,1).
\end{aligned}$$

Then, we can find the pricing kernel of the economy

$$m_{t+1} = \frac{\partial V_t / \partial c_{t+1}}{\partial V_t / \partial c_t}.$$

Now, note that:

$$\frac{\partial V_t}{\partial c_t} = (1-\beta) V_t^{1-\frac{1-\gamma}{\theta}} \upsilon \frac{(c_t^{\upsilon}(1-l_t)^{1-\upsilon})^{\frac{1-\gamma}{\theta}}}{c_t}$$

Author's personal copy

and:

$$\frac{\partial V_t}{\partial c_{t+1}} = \beta V_t^{1-\frac{1-\gamma}{\theta}} \left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{\frac{1}{\theta}-1} \mathbb{E}_t \left(V_{t+1}^{-\gamma}(1-\beta) V_{t+1}^{1-\frac{1-\gamma}{\theta}} \upsilon \frac{(1-\beta)(c_{t+1}^{\upsilon}(1-l_{t+1})^{1-\upsilon})^{\frac{1-\gamma}{\theta}}}{c_{t+1}} \right)$$

where in the last step we use the result regarding $\partial V_t/\partial c_t$ forwarded by one period. Canceling redundant terms, we get:

$$m_{t+1} = \frac{\partial V_t/\partial c_{t+1}}{\partial V_t/\partial c_t} = \beta \left(\frac{c_{t+1}}{c_t}\right)^{\frac{\upsilon(1-\gamma)}{\theta}-1} \left(\frac{1-l_{t+1}}{1-l_t}\right)^{\frac{(1-\gamma)(1-\upsilon)}{\theta}} \left(\frac{V_{t+1}^{1-\gamma}}{\mathbb{E}_t V_{t+1}^{1-\gamma}}\right)^{1-\frac{1}{\theta}}. \tag{2}$$

This equation shows how the pricing kernel is affected by the presence of recursive preferences. If $\theta = 1$, the last term,

$$\left(\frac{V_{t+1}^{1-\gamma}}{\mathbb{E}_t V_{t+1}^{1-\gamma}}\right)^{1-\frac{1}{\theta}} \tag{3}$$

is equal to 1 and we get back the pricing kernel of the standard CRRA case. If $\theta \neq 1$, the pricing kernel is twisted by (3).

We identify the net return on equity with the marginal net return on investment:

$$R_{t+1}^k = \zeta e^{z_{t+1}} k_{t+1}^{\zeta-1} l_{t+1}^{1-\zeta} - \delta$$

with expected return $\mathbb{E}_t[R_{t+1}^k]$.

3. Solution methods

We are interested in comparing different solution methods to approximate the dynamics of the previous model. Since the literature on computational methods is large, it would be cumbersome to review every proposed method. Instead, we select those methods that we find most promising.

Our first method is perturbation (introduced by Judd and Guu, 1992, 1997, and nicely explained in Schmitt-Grohé and Uribe, 2004). Perturbation algorithms build a Taylor series expansion of the agents' decision rules. Often, perturbation methods are very fast and, despite their local nature, highly accurate in a large range of values of the state variables (Aruoba et al., 2006). This means that, in practice, perturbations are the only method that can handle models with dozens of state variables within any reasonable amount of time. Moreover, perturbation often provides insights into the structure of the solution and on the economics of the model. Finally, linearization and log-linearization, the most common solution methods for DSGE models, are particular cases of a perturbation of first order.

We implement a second- and a third-order perturbation of our model. A first-order perturbation is useless for our investigation: the resulting decision rules are certainty equivalent and, therefore, they depend on ψ but not on γ or σ_t. In other words, the first-order decision rules of the model with recursive preferences coincide with the decision rules of the model with CRRA preferences with the same ψ and $\overline{\sigma}$ for *any* value of γ or σ_t. We need to go, at least, to second-order decision rules to have terms that depend on γ or σ_t and, hence, allow recursive preferences or SV to play a role. Since the accuracy of second-order decision rules may not be high enough and, in addition, we want to explore time-varying risk premia, we also compute a third-order perturbation. As we will document below, a third-order perturbation provides enough accuracy without unnecessary complications. Thus, we do not need to go to higher orders.

The second method is a projection algorithm with Chebyshev polynomials (Judd, 1992). Projection algorithms build approximated decision rules that minimize a residual function that measures the distance between the left- and right-hand side of the equilibrium conditions of the model. Projection methods are attractive because they offer a global solution over the whole range of the state space. Their main drawback is that they suffer from an acute curse of dimensionality that makes it challenging to extend it to models with many state variables. Among the many different types of projection methods, Aruoba et al. (2006) show that Chebyshev polynomials are particularly efficient. Other projection methods, such as finite elements or parameterized expectations, tend to perform somewhat worse than Chebyshev polynomials, and therefore, in the interest of space, we do not consider them.

Finally, we compute the model using VFI (Epstein and Zin, 1989, show that a version of the contraction mapping theorem holds in the value function of the problem with recursive preferences). VFI is slow and it suffers as well from the curse of dimensionality, but it is safe, reliable, and well understood. Thus, it is a natural default algorithm for the solution of DSGE models.

3.1. Perturbation

We describe now each of the different methods in more detail. We start by explaining how to use a perturbation approach to solve DSGE models using the value function of the household. We are not the first to explore the perturbation of value function problems. Judd (1998) already presents the idea of perturbing the value function instead of the equilibrium conditions of a model. Unfortunately, he does not elaborate much on the topic. Schmitt-Grohé and Uribe (2005) employ a

Author's personal copy

perturbation approach to find a second-order approximation to the value function that allows them to rank different fiscal and monetary policies in terms of welfare. However, we follow van Binsbergen et al. (2009) in their emphasis on the generality of the approach.[1]

To illustrate the procedure, we limit our exposition to deriving the second-order approximation to the value function and the decision rules of the agents. Higher-order terms are derived analogously, but the algebra becomes too cumbersome to be developed explicitly (in our application, the symbolic algebra is undertaken by `Mathematica`, which automatically generates Fortran 95 code that we can evaluate numerically). Hopefully, our steps will be enough to allow the reader to understand the main thrust of the procedure and obtain higher-order approximations by herself.

First, we rewrite the exogenous processes in terms of a perturbation parameter χ,

$$z_t = \lambda z_{t-1} + \chi e^{\sigma_t} \varepsilon_t$$
$$\sigma_t = (1-\rho)\overline{\sigma} + \rho\sigma_{t-1} + \chi \eta \omega_t.$$

When $\chi = 1$, which is just a normalization, we are dealing with the stochastic version of the model. When $\chi = 0$, we are dealing with the deterministic case with steady state k_{ss}, $z_{ss} = 0$, and $\sigma_{ss} = \overline{\sigma}$. Also, it is convenient for the algebra below to define a vector of states in differences with respect to the steady state, $s_t = (k_t - k_{ss}, z_{t-1}, \varepsilon_t, \sigma_{t-1} - \sigma_{ss}, \omega_t, \chi)$, where s_{it} is the i-th component of this vector at time t for $i \in \{1, \ldots, 6\}$. Then, we can write the social planner's value function, $V(s_t)$, and the decision rules for consumption, $c(s_t)$, investment, $i(s_t)$, capital, $k(s_t)$, and labor, $l(s_t)$, as a function of s_t.

Second, we note that, under differentiability assumptions, the second-order Taylor approximation of the value function around $s_t = \mathbf{0}$ (the vectorial zero) is:

$$V(s_t) \simeq V_{ss} + V_{i,ss} s_t^i + \frac{1}{2} V_{ij,ss} s_t^i s_t^j$$

where:

1. Each term $V_{\ldots,ss}$ is a scalar equal to a derivative of the value function evaluated at $\mathbf{0}$: $V_{ss} \equiv V(\mathbf{0})$, $V_{i,ss} \equiv V_i(\mathbf{0})$ for $i \in \{1, \ldots, 6\}$, and $V_{ij,ss} \equiv V_{ij}(\mathbf{0})$ for $i, j \in \{1, \ldots, 6\}$.
2. We use the tensors $V_{i,ss} s_t^i = \sum_{i=1}^6 V_{i,ss} s_{i,t}$ and $V_{ij,ss} s_t^i s_t^j = \sum_{i=1}^6 \sum_{i=1}^6 V_{ij,ss} s_{i,t} s_{j,t}$, which eliminate the symbol $\sum_{i=1}^6$ when no confusion arises.

We can extend this notation to higher-order derivatives of the value function. This expansion could also be performed around a different point of the state space, such as the mode of the ergodic distribution of the state variables. In Section 5, we discuss this point further.

Fernández-Villaverde et al. (2010) show that many of these terms $V_{\ldots,ss}$ are zero (for instance, those implied by certainty equivalence in the first-order component). More directly related to this paper, van Binsbergen et al. (2009) demonstrate that γ does not affect the values of any of the coefficients except $V_{66,ss}$ and also that $V_{66,ss} \neq 0$. This result is intuitive, since the value function of a risk-averse agent is in general affected by uncertainty and we want to have an approximation with terms that capture this effect and allow for the appropriate welfare ranking of decision rules. Indeed, $V_{66,ss}$ has a straightforward interpretation. At the deterministic steady state with $\chi = 1$ (that is, even if we are in the stochastic economy, we just happen to be exactly at the steady-state values of all the other states), we have:

$$V(0,0,0,0,0,1) \simeq V_{ss} + \frac{1}{2} V_{66,ss}.$$

Hence $\frac{1}{2} V_{66,ss}$ is a measure of the welfare cost of the business cycle, that is, of how much utility changes when the variance of the productivity shocks is at steady-state value σ_{ss} instead of zero (note that this quantity is not necessarily negative). This term is an accurate evaluation of the third order of the welfare cost of business cycle fluctuations because all of the third-order terms in the approximation of the value function either have coefficient values of zero or drop when evaluated at the deterministic steady state.

This cost of the business cycle can easily be transformed into consumption equivalent units. We can compute the percentage decrease in consumption τ that will make the household indifferent between consuming $(1-\tau)c_{ss}$ units per period with certainty or c_t units with uncertainty. To do so, note that the steady-state value function is just $V_{ss} = c_{ss}^{\upsilon}(1-l_{ss})^{1-\upsilon}$, which implies that:

[1] The perturbation method is related to Benigno and Woodford (2006) and Hansen and Sargent (1995). Benigno and Woodford present a linear-quadratic (LQ) approximation to solve optimal policy problems when the constraints of the problem are non-linear (see also Levine et al., 2007). This allows them to find the correct local welfare ranking of different policies. Our perturbation can also deal with non-linear constraints and obtains the correct local approximation to welfare and policies, but with the advantage that it is easily generalizable to higher-order approximations. Hansen and Sargent (1995) modify the LQ problem to adjust for risk. In that way, they can handle some versions of recursive utilities. Hansen and Sargent's method, however, requires imposing a tight functional form for future utility and to surrender the assumption that risk-adjusted utility is separable across states of the world. Perturbation does not suffer from those limitations.

Author's personal copy

$$c_{ss}^{\upsilon}(1-l_{ss})^{1-\upsilon}+\frac{1}{2}V_{66,ss}=\big((1-\tau)c_{ss}\big)^{\upsilon}(1-l_{ss})^{1-\upsilon}$$

or:

$$V_{ss}+\frac{1}{2}V_{66,ss}=(1-\tau)^{\upsilon}V_{ss}.$$

Then:

$$\tau=1-\left[1+\frac{1}{2}\frac{V_{66,ss}}{V_{ss}}\right]^{\frac{1}{\upsilon}}.$$

We are perturbing the value function in levels of the variables. However, there is nothing special about levels and we could have done the same in logs (a common practice when linearizing DSGE models) or in any other function of the states. These changes of variables may improve the performance of perturbation (Fernández-Villaverde and Rubio-Ramírez, 2006). By doing the perturbation in levels, we are picking the most conservative case for perturbation. Since one of the conclusions that we will reach from our numerical results is that perturbation works surprisingly well in terms of accuracy, that conclusion will only be reinforced by an appropriate change of variables.[2]

The decision rules can be expanded in the same way. For example, the second-order approximation of the decision rule for consumption is, under differentiability assumptions:

$$c(s_t)\simeq c_{ss}+c_{i,ss}s_t^i+\frac{1}{2}c_{ij,ss}s_t^i s_t^j$$

where we have followed the same derivative and tensor notation as before.

As with the approximation of the value function, van Binsbergen et al. (2009) show that γ does not affect the values of any of the coefficients except $c_{66,ss}$. This term is a constant that captures precautionary behavior caused by risk. This observation tells us two facts. First, a linear approximation to the decision rule does not depend on γ (it is certainty equivalent), and therefore, if we are interested in recursive preferences, we need to go at least to a second-order approximation. Second, given some fixed parameter values, the difference between the second-order approximation to the decision rules of a model with CRRA preferences and a model with recursive preferences is a constant. This constant generates a second, indirect effect, because it changes the ergodic distribution of the state variables and, hence, the points where we evaluate the decision rules along the equilibrium path. These arguments demonstrate how perturbation methods can provide analytic insights beyond computational advantages and help in understanding the numerical results in Tallarini (2000). In the third-order approximation, all of the terms on functions of χ^2 depend on γ.

Following the same steps, we can derive the decision rules for labor, investment, and capital. In addition we have functions that give us the evolution of other variables of interest, such as the pricing kernel or the risk-free gross interest rate R_t^f. All of these functions have the same structure and properties regarding γ as the decision rule for consumption. In the case of functions pricing assets, the second-order approximation generates a constant risk premium, while the third-order approximation creates a time-varying risk premium.

Once we have reached this point, there are two paths we can follow to solve for the coefficients of the perturbation. The first procedure is to write down the equilibrium conditions of the model plus the definition of the value function. Then, we take successive derivatives in this augmented set of equilibrium conditions and solve for the unknown coefficients. This approach, which we call equilibrium conditions perturbation (ECP), gets us, after n iterations, the n-th-order approximation to the value function and to the decision rules.

A second procedure is to take derivatives of the value function with respect to states and controls and use those derivatives to find the unknown coefficient. This approach, which we call value function perturbation (VFP), delivers after $(n+1)$ steps, the $(n+1)$-th-order approximation to the value function and the n-th-order approximation to the decision rules.[3] Loosely speaking, ECP undertakes the first step of VFP by hand by forcing the researcher to derive the equilibrium conditions.

The ECP approach is simpler but it relies on our ability to find equilibrium conditions that do not depend on derivatives of the value function. Otherwise, we need to modify the equilibrium conditions to include the definitions of the derivatives of the value function. Even if this is possible to do (and not particularly difficult), it amounts to solving a problem that is equivalent to VFP. This observation is important because it is easy to postulate models that have equilibrium conditions where we cannot get rid of all the derivatives of the value function (for example, in problems of optimal policy design). ECP is also faster from a computational perspective. However, VFP is only trivially more involved because finding the $(n+1)$-th-order approximation to the value function on top of the n-th-order approximation requires nearly no additional effort.

[2] This comment begets the question, nevertheless, of why we did not perform a perturbation in logs, since many economists will find it more natural than using levels. Our experience with the CRRA utility case (Aruoba et al., 2006) and some computations with recursive preferences not included in the paper suggest that a perturbation in logs does slightly worse than a perturbation in levels.

[3] The classical strategy of finding a quadratic approximation of the utility function to derive a linear decision rule is a second-order example of VFP (Anderson et al., 1996). A standard linearization of the equilibrium conditions of a DSGE model when we add the value function to those equilibrium conditions is a simple case of ECP. This is done, for instance, in Schmitt-Grohé and Uribe (2005).

Author's personal copy

The algorithm presented here is based on the system of equilibrium equations derived using the ECP. In Appendix A, we derive a system of equations using the VFP. We take the first-order conditions of the social planner. First, with respect to consumption today:

$$\frac{\partial V_t}{\partial c_t} - \mu_t = 0$$

where μ_t is the Lagrangian multiplier associated with the resource constraint. Second, with respect to capital:

$$-\mu_t + \mathbb{E}_t \mu_{t+1}\left(\zeta e^{z_{t+1}} k_{t+1}^{\zeta-1} l_{t+1}^{1-\zeta} + 1 - \delta\right) = 0.$$

Third, with respect to labor:

$$\frac{1-\upsilon}{\upsilon}\frac{c_t}{(1-l_t)} = (1-\zeta)e^{z_t} k_t^{\zeta} l_t^{-\zeta}.$$

Then, we have $\mathbb{E}_t m_{t+1}(\zeta e^{z_t} k_{t+1}^{\zeta-1} l_{t+1}^{1-\zeta} + 1 - \delta) = 1$ where m_{t+1} was derived in Eq. (2). Note that, as explained above, the derivatives of the value function in (2) are eliminated.

Once we substitute for the pricing kernel, the augmented equilibrium conditions are:

$$V_t - \left[(1-\beta)\left(c_t^{\upsilon}(1-l_t)^{1-\upsilon}\right)^{\frac{1-\gamma}{\theta}} + \beta\left(\mathbb{E}_t V^{1-\gamma}(k_{t+1}, z_{t+1})\right)^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}} = 0$$

$$\mathbb{E}_t\left[\beta\left(\frac{c_{t+1}}{c_t}\right)^{\frac{1-\gamma}{\theta}-1}\left(\frac{V_{t+1}^{1-\gamma}}{\mathbb{E}_t V_{t+1}^{1-\gamma}}\right)^{1-\frac{1}{\theta}}\left(\zeta e^{z_{t+1}} k_{t+1}^{\zeta-1} l_{t+1}^{1-\zeta} + 1 - \delta\right)\right] - 1 = 0$$

$$\frac{1-\upsilon}{\upsilon}\frac{c_t}{(1-l_t)} = (1-\zeta)e^{z_t} k_t^{\zeta} l_t^{-\zeta} = 0$$

$$\mathbb{E}_t \beta\left(\frac{c_{t+1}}{c_t}\right)^{\frac{1-\gamma}{\theta}-1}\left(\frac{V_{t+1}^{1-\gamma}}{\mathbb{E}_t V_{t+1}^{1-\gamma}}\right)^{1-\frac{1}{\theta}} R_t^f - 1 = 0$$

$$c_t + i_t - e^{z_t} k_t^{\zeta} l_t^{1-\zeta} = 0$$

$$k_{t+1} - i_t - (1-\delta)k_t = 0$$

plus the law of motion for productivity and volatility. Note that all the endogenous variables are functions of the states and that we drop the max operator in front of the value function because we are already evaluating it at the optimum. Thus, a more compact notation for the previous equilibrium conditions as a function of the states is:

$$F(\mathbf{0}) = \mathbf{0}$$

where $F : \mathbb{R}^6 \to \mathbb{R}^8$.

In steady state, $m_{ss} = \beta$ and the set of equilibrium conditions simplifies to:

$$V_{ss} = c_{ss}^{\upsilon}(1-l_{ss})^{1-\upsilon}$$

$$\left(\zeta k_{ss}^{\zeta-1} l_{ss}^{1-\zeta} + 1 - \delta\right) = 1/\beta$$

$$\frac{1-\upsilon}{\upsilon}\frac{c_{ss}}{(1-l_{ss})} = (1-\zeta)k_{ss}^{\zeta} l_{ss}^{-\zeta}$$

$$R_{ss}^f = 1/\beta$$

$$c_{ss} + i_{ss} = k_{ss}^{\zeta} l_{ss}^{1-\zeta}$$

$$i_{ss} = \delta k_{ss}$$

a system of 6 equations on 6 unknowns, V_{ss}, c_{ss}, k_{ss}, i_{ss}, l_{ss}, and R_{ss}^f that can be easily solved (see Appendix A for the derivations). This steady state is identical to the steady state of the real business cycle model with a standard CRRA utility function and no SV.

To find the first-order approximation to the value function and the decision rules, we take first derivatives of the function F with respect to the states s_t and evaluate them at $\mathbf{0}$:

$$F_i(\mathbf{0}) = \mathbf{0} \quad \text{for } i \in \{1, \ldots, 6\}.$$

This step gives us 48 different first derivatives (8 equilibrium conditions times the 6 variables of F). Since each dimension of F is equal to zero for all possible values of s_t, their derivatives must also be equal to zero. Therefore, once we substitute

Author's personal copy

the steady-state values and forget about the exogenous processes (which we do not need to solve for), we have a quadratic system of 36 equations on 36 unknowns: $V_{i,ss}$, $c_{i,ss}$, $i_{i,ss}$, $k_{i,ss}$, $l_{i,ss}$, and $R^f_{i,ss}$ for $i \in \{1, \dots, 6\}$. One of the solutions is an unstable root of the system that violates the transversality condition of the problem and we eliminate it. Thus, we keep the solution that implies stability.

To find the second-order approximation, we take derivatives on the first derivatives of the function F, again with respect to the states and the perturbation parameter:

$$F_{ij}(\mathbf{0}) = \mathbf{0} \quad \text{for } i, j \in \{1, \dots, 6\}.$$

This step gives us a new system of equations. Then, we plug in the terms that we already know from the steady state and from the first-order approximation and we get that the only unknowns left are the second-order terms of the value function and other functions of interest. Quite conveniently, this system of equations is linear and it can be solved quickly. Repeating these steps (taking higher-order derivatives, plugging in the terms already known, and solving for the remaining unknowns), we can get any arbitrary order approximation. For simplicity, and since we checked that we were already obtaining a high accuracy, we decided to stop at a third-order approximation (we are particularly interested in applying the perturbation for estimation purposes and we want to document how a third-order approximation is accurate enough for many problems without spending too much time deriving higher-order terms).

3.2. Projection

Projection methods take basis functions to build an approximated value function and decision rules that minimize a residual function defined by the augmented equilibrium conditions of the model. There are two popular methods for choosing basis functions: finite elements and the spectral method. We will present only the spectral method for several reasons: first, in the neoclassical growth model the decision rules and value function are smooth and spectral methods provide an excellent approximation. Second, spectral methods allow us to use a large number of basis functions, with the consequent high accuracy. Third, spectral methods are easier to implement. Their main drawback is that since they approximate the solution with a spectral basis, if the decision rules display a rapidly changing local behavior or kinks, it may be difficult for this scheme to capture those local properties.

Our target is to solve the decision rule for labor and the value function $\{l_t, V_t\}$ from:

$$\mathcal{H}(l_t, V_t) = \begin{bmatrix} u_{c,t} - \beta\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{\frac{1}{\theta}-1} \mathbb{E}_t\left[V_{t+1}^{\frac{(1-\gamma)(\theta-1)}{\theta}} u_{c,t+1}\left(\zeta e^{z_{t+1}} k_{t+1}^{\zeta-1} l_{t+1}^{1-\zeta} + 1 - \delta\right)\right] \\ V_t - \left[(1-\beta)\left(c_t^{\upsilon}\left(1 - l_t^{\upsilon}\right)\right)^{\frac{1-\gamma}{\theta}} + \beta \mathbb{E}_t\left(V_{t+1}^{1-\gamma}\right)^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}} \end{bmatrix} = \mathbf{0}$$

where, to save on notation, we define $V_t = V(k_t, z_t, \sigma_t)$ and:

$$u_{c,t} = \frac{1-\gamma}{\theta} \upsilon \frac{(c_t^{\upsilon}(1-l_t)^{1-\upsilon})^{\frac{1-\gamma}{\theta}}}{c_t}.$$

Then, from the static condition

$$c_t = \frac{\upsilon}{1-\upsilon}(1-\zeta) e^{z_t} k_t^{\zeta} l_t^{-\zeta} (1 - l_t)$$

and the resource constraint, we can find c_t and k_{t+1}.

Spectral methods solve this problem by specifying the decision rule for labor and the value function $\{l_t, V_t\}$ as linear combinations of weighted basis functions:

$$l(k_t, z_j, \sigma_m; \rho) = \sum_i \rho^l_{ijm} \psi_i(k_t)$$

$$V(k_t, z_j, \sigma_m; \rho) = \sum_i \rho^V_{ijm} \psi_i(k_t)$$

where $\{\psi_i(k)\}_{i=1,\dots,n_k}$ are the n_k basis functions that we will use for our approximation along the capital dimension and $\rho = \{\rho^l_{ijm}, \rho^V_{ijm}\}_{i=1,\dots,n_k;\ j=1,\dots,J;\ m=1,\dots,M}$ are unknown coefficients to be determined. In this expression, we have discretized the stochastic processes σ_t for volatility and z_t for productivity using Tauchen's (1986) method as follows. First, we have a grid of M points $G_\sigma = \{e^{\sigma_1}, e^{\sigma 2}, \dots, e^{\sigma_M}\}$ for σ_t and a transition matrix Π^M with generic element $\pi^M_{i,j} = Prob(e^{\sigma_{t+1}} = e^{\sigma_j} | e^{\sigma_t} = e^{\sigma_i})$. The grid covers 3 standard deviations of the process in each direction. Then, for each M point, we find a grid with J points $G^m_z = \{z^m_1, z^m_2, \dots, z^m_J\}$ for z_t and transition matrices $\Pi^{J,m}$ with generic element $\pi^{J,m}_{i,j} = Prob(z^m_{t+1} = z^m_j | z^m_t = z^m_i)$. Again, and conditional on e^{σ_m}, the grid covers 3 standard deviations in each direction. Values for the decision rule outside the grids G_σ and G^m_z can be approximated by interpolation. We make the grids for z_t depend on the level of volatility m to adapt the accuracy of Tauchen's procedure to each conditional variance (although this forces us to interpolate when we switch variances). Since we set $J = 25$ and $M = 5$, the approximation is quite accurate along the productivity axis (we explored other choices of J and M to be sure that our choice was sensible).

Author's personal copy

A common choice for the basis functions are Chebyshev polynomials because of their flexibility and convenience. Since their domain is $[-1, 1]$, we need to bound capital to the set $[\underline{k}, \bar{k}]$, where $\underline{k}$ ($\bar{k}$) is chosen sufficiently low (high) to bind only with extremely low probability, and define a linear map from those bounds into $[-1, 1]$. Then, we set $\psi_i(k_t) = \tilde{\psi}_i(\phi_k(k_t))$ where $\tilde{\psi}_i(\cdot)$ are Chebyshev polynomials and $\phi_k(k_t)$ is the linear mapping from $[\underline{k}, \bar{k}]$ to $[-1, 1]$.

By plugging $l(k_t, z_j, \sigma_m; \rho)$ and $V(k_t, z_j, \sigma_m; \rho)$ into $\mathcal{H}(l_t, V_t)$, we find the residual function:

$$\mathcal{R}(k_t, z_j, \sigma_m; \rho) = \mathcal{H}\big(l(k_t, z_j, \sigma_m; \rho), V(k_t, z_j, \sigma_m; \rho)\big).$$

We determine the coefficients ρ to get the residual function as close to $\mathbf{0}$ as possible. However, to do so, we need to choose a weight of the residual function over the space (k_t, z_j, σ_m). A collocation point criterion delivers the best trade-off between speed and accuracy (Fornberg, 1998) by making the residual function exactly equal to zero in $\{k_i\}_{i=1}^{n_k}$ roots of the n_k-th-order Chebyshev polynomial and in the Tauchen points (also, by the Chebyshev interpolation theorem, if an approximating function is exact at the roots of the n_k-th-order Chebyshev polynomial, then, as $n_k \to \infty$, the approximation error becomes arbitrarily small). Therefore, we just need to solve the following system of $n_k \times J \times M \times 2$ equations:

$$\mathcal{R}(k_i, z_j, \sigma_m; \rho) = \mathbf{0} \quad \text{for any } i, j, m \text{ collocation points}$$

on $n_k \times J \times M \times 2$ unknowns ρ. We solve this system with a Newton method and an iteration based on the increment of the number of basis functions. First, we solve a system with only three collocation points for capital and 125 ($125 = 25 \times 5$) points for technology. Then, we use that solution as a guess for a system with more collocation points for capital (with the new coefficients being guessed to be equal to zero) and iterate on the procedure. We stop the iteration when we have 11 polynomials in the capital dimension (therefore, in the last step we solve for $2750 = 11 \times 25 \times 5 \times 2$ coefficients). The iteration is needed because otherwise the residual function is too cumbersome to allow for direct solution of the 2750 final coefficients.

3.3. Value function iteration

Our final solution method is VFI. Since the dynamic algorithm is well known, our presentation is most brief. Consider the following Bellman operator:

$$\begin{aligned}
TV(k_t, z_t, \sigma_t) &= \max_{c_t, l_t} \Big[(1-\beta)\big(c_t^{\upsilon}(1-l_t)^{1-\upsilon}\big)^{\frac{1-\gamma}{\theta}} + \beta\big(\mathbb{E}_t V^{1-\gamma}(k_{t+1}, z_{t+1}, \sigma_{t+1})\big)^{\frac{1}{\theta}}\Big]^{\frac{\theta}{1-\gamma}} \\
\text{s.t.} \quad & c_t + k_{t+1} = e^{z_t} k_t^{\zeta} l_t^{1-\zeta} + (1-\delta)k_t \\
& z_t = \lambda z_{t-1} + e^{\sigma_t}\varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0,1) \\
& \sigma_t = (1-\rho)\bar{\sigma} + \rho\sigma_{t-1} + \eta\omega_t, \quad \omega_t \sim \mathcal{N}(0,1).
\end{aligned}$$

To solve for this Bellman operator, we define a grid on capital, $G_k = \{k_1, k_2, \ldots, k_M\}$, a grid on volatility and on the productivity level. The grid on capital is just a uniform distribution of points over the capital dimension. As we did for projection, we set a grid $G_\sigma = \{e^{\sigma_1}, e^{\sigma 2}, \ldots, e^{\sigma_M}\}$ for σ_t and a transition matrix Π^M for volatility and M grids $G_z^m = \{z_1^m, z_2^m, \ldots, z_J^m\}$ for z_t and transition matrices $\Pi^{J,m}$ using Tauchen's (1986) procedure. The algorithm to iterate on the value function for this grid is:

1. Set $n = 0$ and $V^0(k_t, z_t, \sigma_t) = c_{ss}^{\upsilon}(1-l_{ss})^{1-\upsilon}$ for all $k_t \in G_k$ and all $z_t \in G_z$.
2. For every $\{k_t, z_t, \sigma_t\}$, use the Newton method to find c_t^*, l_t^*, k_{t+1}^* that solve:

$$\begin{aligned}
c_t &= \frac{\upsilon}{1-\upsilon}(1-\zeta)e^{z_t}k_t^{\zeta}l_t^{-\zeta}(1-l_t) \\
(1-\beta)\upsilon\frac{(c_t^{\upsilon}(1-l_t)^{1-\upsilon})^{\frac{1-\gamma}{\theta}}}{c_t} &= \beta\big(\mathbb{E}_t\big(V_{t+1}^n\big)^{1-\gamma}\big)^{\frac{1}{\theta}-1}\mathbb{E}_t\big[\big(V_{t+1}^n\big)^{-\gamma}V_{1,t+1}^n\big] \\
c_t + k_{t+1} &= e^{z_t}k_t^{\zeta}l_t^{1-\zeta} + (1-\delta)k_t.
\end{aligned}$$

3. Construct V^{n+1} from the Bellman equation:

$$V^{n+1} = \Big((1-\beta)\big(c_t^{*\upsilon}\big(1-l_t^*\big)^{1-\upsilon}\big)^{\frac{1-\gamma}{\theta}} + \beta\big(\mathbb{E}_t\big(V\big(k_{t+1}^*, z_{t+1}, \sigma_{t+1}\big)^{1-\gamma}\big)\big)^{\frac{1}{\theta}}\Big)^{\frac{\theta}{1-\gamma}}.$$

4. If $\frac{|V^{n+1}-V^n|}{|V^n|} \geqslant 1.0e^{-7}$, then $n = n+1$ and go to 2. Otherwise, stop.

To accelerate convergence and give VFI a fair chance, we implement a multigrid scheme as described by Chow and Tsitsiklis (1991). We start by iterating on a small grid. Then, after convergence, we add more points to the grid and recompute

Author's personal copy

Table 1
Calibrated parameters.

Parameter	β	υ	ζ	δ	λ	$\log\overline{\sigma}$	ρ	η
Value	0.991	0.357	0.3	0.0196	0.95	0.007	0.9	0.06

the Bellman operator using the previously found value function as an initial guess (with linear interpolation to fill the unknown values in the new grid points). Since the previous value function is an excellent grid, we quickly converge in the new grid. Repeating these steps several times, we move from an initial 23,000-point grid into a final one with 375,000 points (3000 points for capital, 25 for productivity, and 5 for volatility).

4. Calibration

We now select a benchmark calibration for our numerical computations. We follow the literature as closely as possible and select parameter values to match, in the steady state, some basic observations of the U.S. economy (as we will see below, for the benchmark calibration, the means of the ergodic distribution and the steady-state values are nearly identical). We set the discount factor $\beta = 0.991$ to generate an annual interest rate of around 3.6 percent. We set the parameter that governs labor supply, $\theta = 0.357$, to get the representative household to work one-third of its time. The elasticity of output to capital, $\zeta = 0.3$, matches the labor share of national income. A value of the depreciation rate $\delta = 0.0196$ matches the ratio of investment-output. Finally, $\lambda = 0.95$ and $\log\overline{\sigma} = 0.007$ are standard values for the stochastic properties of the Solow residual. For the SV process, we pick $\rho = 0.9$ and $\eta = 0.06$, to match the persistence and standard deviation of the heteroskedastic component of the Solow residual during the last 5 decades. (See Table 1.)

Since we want to explore the dynamics of the model for a range of values that encompasses all the estimates from the literature, we select four values for the parameter that controls risk aversion, γ, 2, 5, 10, and 40, and two values for EIS ψ, 0.5, and 1.5, which bracket most of the values used in the literature (although many authors prefer smaller values for ψ, we found that the simulation results for smaller ψ's do not change much from the case when $\psi = 0.5$). We then compute the model for all eight combinations of values of γ and ψ, that is $\{2, 0.5\}$, $\{5, 0.5\}$, $\{10, 0.5\}$, and so on. When $\psi = 0.5$ and $\gamma = 2$, we are back in the standard CRRA case. However, in the interest of space, we will report only a limited subset of results that we find are the most interesting ones.

We pick as the benchmark case the calibration $\{\gamma, \psi, \log\overline{\sigma}, \eta\} = \{5, 0.5, 0.007, 0.06\}$. These values reflect an EIS centered around the median of the estimates in the literature, a reasonably high level of risk aversion, and the observed volatility of productivity shocks. To check robustness, we increase, in the extreme case, the risk aversion, the average standard deviation of the productivity shock, and the standard deviation of the innovations to volatility to $\{\gamma, \psi, \log\overline{\sigma}, \eta\} = \{40, 0.5, 0.021, 0.1\}$. This case combines levels of risk aversion that are in the upper bound of all estimates in the literature with huge productivity shocks. Therefore, it pushes all solution methods to their limits, in particular, making life hard for perturbation since the interaction of the large precautionary behavior induced by γ and large shocks will move the economy far away from the deterministic steady state. We leave the discussion of the effects of $\psi = 1.5$ for the robustness analysis at the end of the next section.

5. Numerical results

In this section we report our numerical findings. First, we present and discuss the computed decision rules. Second, we show the results of simulating the model. Third, we report the Euler equation errors. Fourth, we discuss the effects of changing the EIS and the perturbation point. Finally, we discuss implementation and computing time.

5.1. Decision rules

Fig. 1 plots the decision rules of the household for consumption and labor in the benchmark case. In the top two panels, we plot the decision rules along the capital axis when $z_t = 0$ and $\sigma_t = \overline{\sigma}$. The capital interval is centered on the steady-state level of capital of 9.54 with a width of ±40%, [5.72,13.36]. This interval is big enough to encompass all the simulations. We also plot, with two vertical lines, the 10 and 90 percentile of the ergodic distribution of capital.[4] Since all methods provide nearly indistinguishable answers, we observe only one line in both panels. It is possible to appreciate miniscule differences in labor supply between second-order perturbation and the other methods only when capital is far from its steady-state level. Monotonicity of the decision rules is preserved by all methods.[5] We must be cautious, however, mapping differences in choices into differences in utility. The Euler error function below provides a better view of the welfare consequences of different approximations.

[4] There is the technical consideration of which ergodic distribution to use for this task, since this is an object that can only be found by simulation. All over the paper, we use the ergodic simulation generated by VFI. We checked that the results are robust to using the ergodic distributions coming from the other methods.

[5] Similar figures could be plotted for other values of z_t and σ_t. We omit them because of space considerations.

Author's personal copy

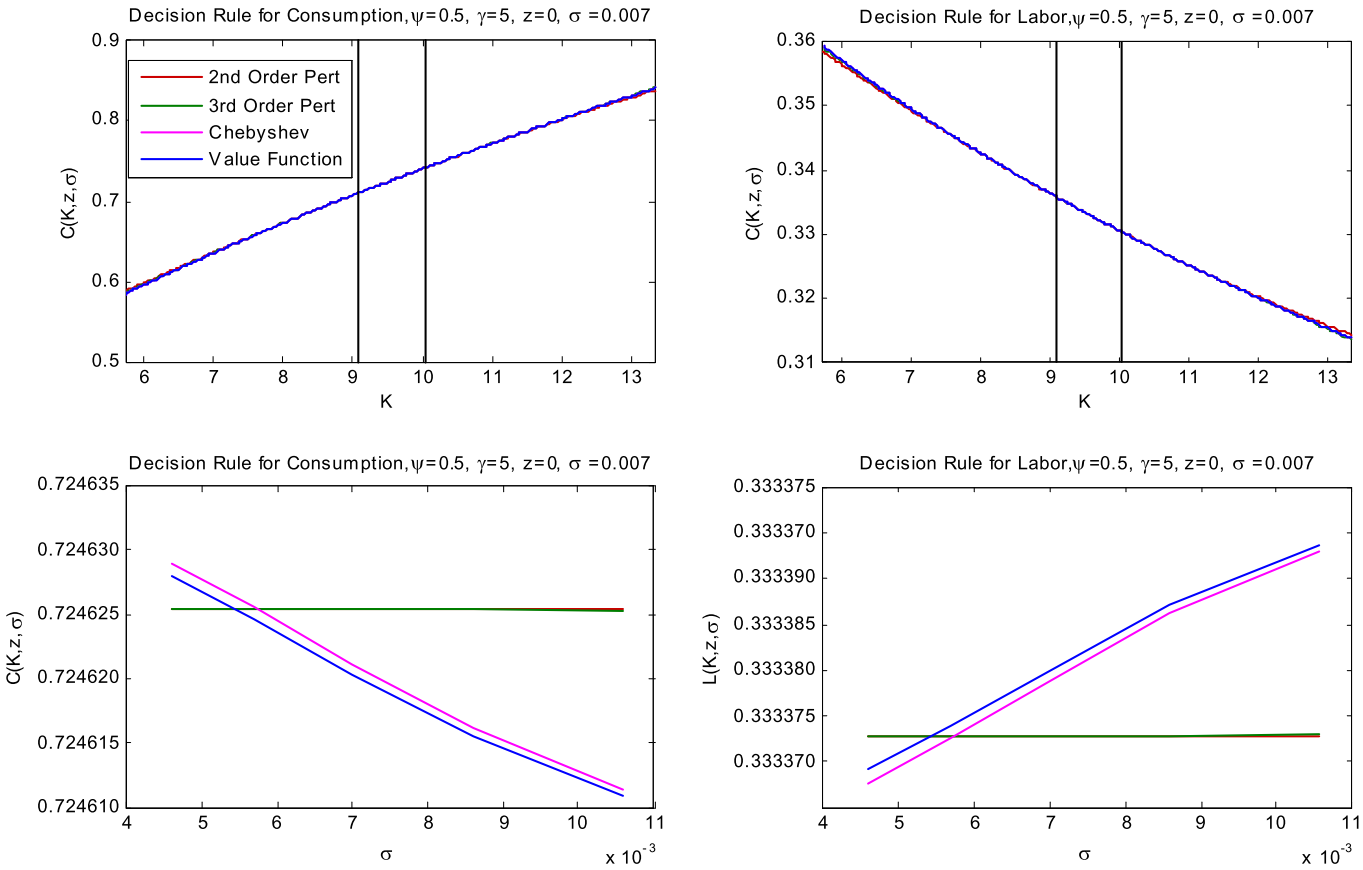

Fig. 1. Decision rules, benchmark calibration.

In the bottom two panels, we plot the decision rules as a function of σ_t when $k_t = k_{ss}$ and $z_t = 0$. Both VFI and Chebyshev polynomials capture well precautionary behavior: the household consumes less and works more when σ_t is higher. In comparison, it would seem that perturbations do a much poorer job at capturing this behavior. However, this is not the right interpretation. In a second-order perturbation, the effect of σ_t only appears in a term where σ_t interacts with z_t. Since here we are plotting a cut of the decision rule while keeping $z_t = 0$, that term drops out and the decision rule that comes from a second-order approximation is a flat line. In the third-order perturbation, σ_t appears by itself, but raised to its cube. Hence, the slope of the decision rule is negligible. As we will see below, the interaction effects of σ_t with other variables, which are hard to visualize in two-dimensional graph, are all that we need to deliver a satisfactory overall performance.

We see bigger differences in the decision rules as we increase the risk aversion and the variance of innovations. Fig. 2 plots the decision rules for the extreme calibration following the same convention than before. In this figure, we change the capital interval where we compute the top decision rules to $[3, 32]$ (roughly $1/3$ and 3 times the steady-state capital) because, owing to the high variance of the calibration, the equilibrium paths fluctuate through much wider ranges of capital.

We highlight several results. First, all the methods deliver similar results in our original capital interval for the benchmark calibration. Second, as we go far away from the steady state, VFI and the Chebyshev polynomial still overlap with each other (and, as shown by our Euler error computations below, we can roughly take them as the "exact" solution), but second- and third-order approximations start to deviate. Third, the decision rule for consumption approximated by the third-order perturbation changes from concavity into convexity for values of capital bigger than 15. This phenomenon (also documented in Aruoba et al., 2006) is due to the poor performance of local approximation when we move too far away from the expansion point and the polynomials begin to behave wildly. Numerically, this issue is irrelevant because the problematic region is visited with nearly zero probability.

5.2. Simulations

Applied economists often characterize the behavior of the model through statistics from simulated paths of the economy. We simulate the model, starting from the deterministic steady state, for 10,000 periods, using the decision rules for each of the eight combinations of risk aversion and EIS discussed above. To make the comparison meaningful, the shocks are common across all paths. We discard the first 1000 periods as a burn-in to eliminate the transition from the deterministic steady state of the model to the middle regions of the ergodic distribution of capital. This is usually achieved in many fewer

Author's personal copy

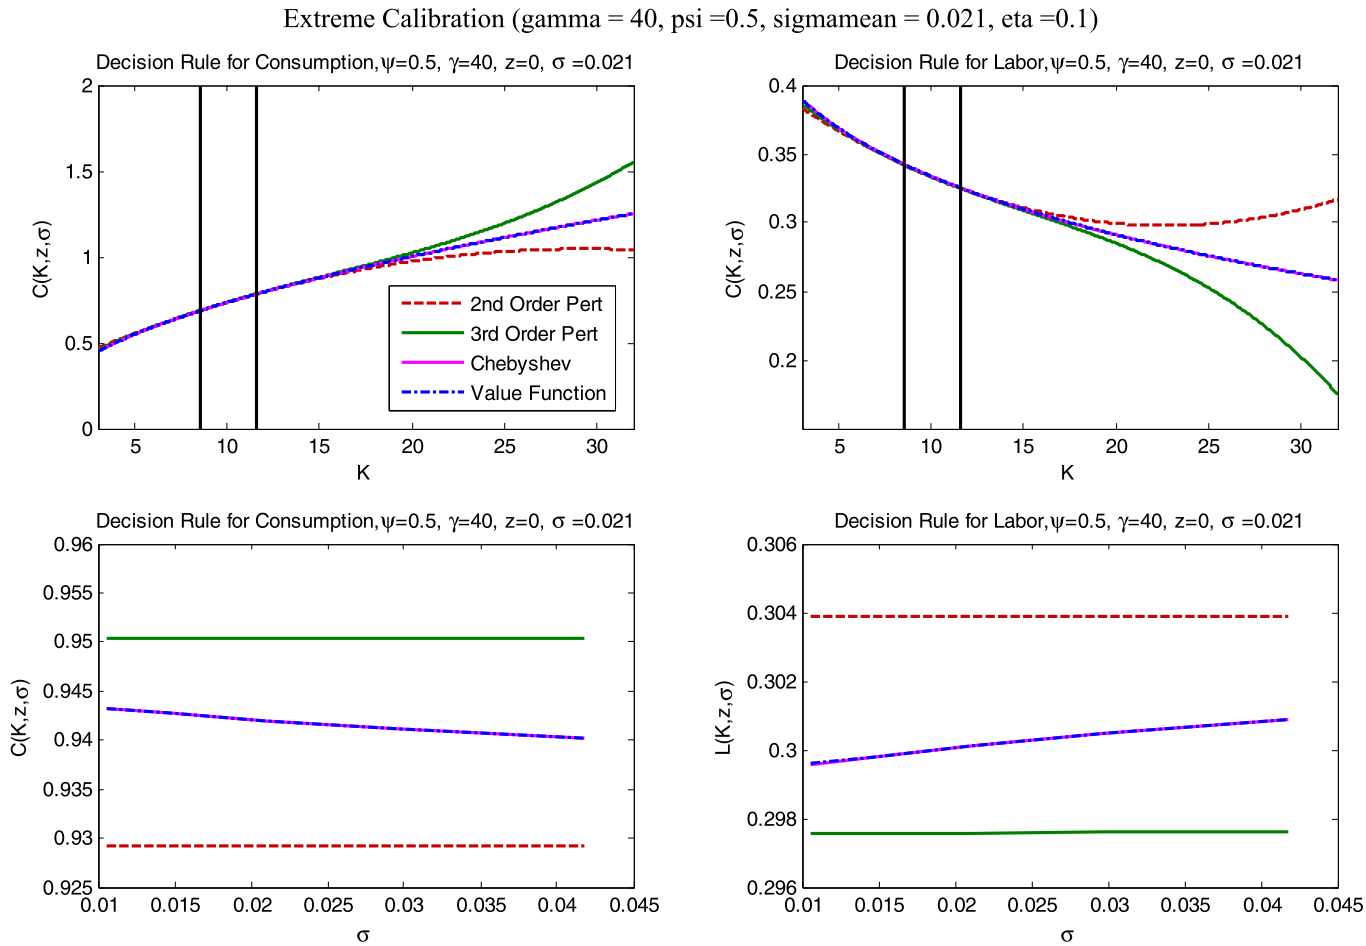

Fig. 2. Decision rules, extreme calibration.

periods than the ones in our burn-in, but we want to be conservative in our results. The remaining observations constitute a sample from the ergodic distribution of the economy.

For the benchmark calibration, the simulations from all of the solution methods generate almost identical equilibrium paths (and therefore we do not report them). We focus instead on the densities of the endogenous variables as shown in Fig. 3 (remember that volatility and productivity are identical across the different solution methods). Given the low risk aversion and SV of the productivity shocks, all densities are roughly centered around the deterministic steady-state value of the variable. For example, the mean of the distribution of capital is only 0.2 percent higher than the deterministic value. Also, capital is nearly always between 8.5 and 10.5. This range will be important below to judge the accuracy of our approximations.

Table 2 reports business cycle statistics and, because DSGE models with recursive preferences and SV are often used for asset pricing, the average and variance of the (quarterly) risk-free rate and return on capital. Again, we see that nearly all values are the same, a simple consequence of the similarity of the decision rules.

The welfare cost of the business cycle is reported in Table 3 in consumption equivalent terms. The computed costs are actually negative. Besides the Jensen's effect on average productivity, this is also due to the fact that when we have leisure in the utility function, the indirect utility function may be convex in input prices (agents change their behavior over time by a large amount to take advantage of changing productivity). Cho and Cooley (2000) present a similar example. Welfare costs are comparable across methods. Remember that the welfare cost of the business cycle for the second- and third-order perturbations is the same because the third-order terms all drop or are zero when evaluated at the steady state.

When we move to the extreme calibration, we see more differences. Fig. 4 plots the histograms of the simulated series for each solution method. Looking at quantities, the histograms of consumption, output, and labor are the same across all of the methods. The ergodic distribution of capital puts nearly all the mass between values of 6 and 15. This considerable move to the right in comparison with Fig. 3 is due to the effect of precautionary behavior in the presence of high risk aversion, large productivity shocks, and high SV. Capital also visits low values of capital more than in the benchmark calibration because of large, persistent productivity shocks. In any case, the translation is more pronounced to the right than to the left.

Table 4 reports business cycle statistics. Differences across methods are minor in terms of means (note that the mean of the risk-free rate is lower than in the benchmark calibration because of the extra accumulation of capital induced by

Author's personal copy

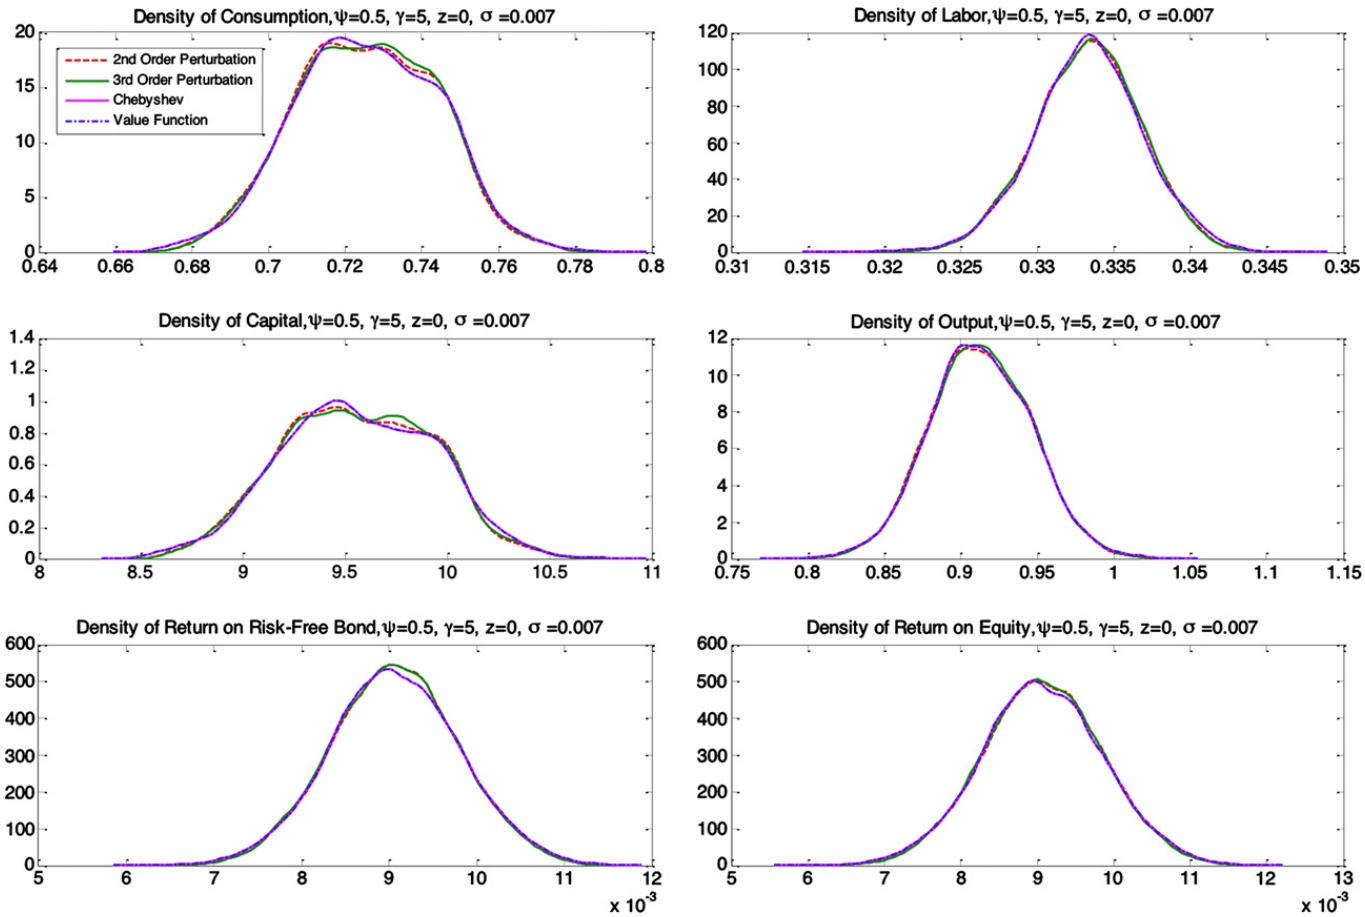

Fig. 3. Densities, benchmark calibration.

Table 2
Business cycle statistics – benchmark calibration.

	c	y	i	R^f (%)	R^k (%)
	Mean				
Second-order perturbation	0.7253	0.9128	0.1873	0.9070	0.9078
Third-order perturbation	0.7257	0.9133	0.1875	0.9062	0.9069
Chebyshev polynomial	0.7256	0.9130	0.1875	0.9063	0.9066
Value function iteration	0.7256	0.9130	0.1875	0.9063	0.9066
	Variance (%)				
Second-order perturbation	0.0331	0.1084	0.0293	0.0001	0.0001
Third-order perturbation	0.0330	0.1079	0.0288	0.0001	0.0001
Chebyshev polynomial	0.0347	0.1117	0.0313	0.0001	0.0001
Value function iteration	0.0347	0.1117	0.0313	0.0001	0.0001

Table 3
Welfare costs of business cycle – benchmark calibration.

2nd-order pert.	3rd-order pert.	Chebyshev	Value function
−2.0864e(−5)	−2.0864e(−5)	−3.2849e(−5)	−3.2849e(−5)

precautionary behavior). In terms of variances, the second-order perturbation produces less volatility than all other methods. This suggests that a second-order perturbation may not be good enough if we face high variance of the shocks and/or high risk aversion. A third-order perturbation, in comparison, eliminates most of the differences and delivers nearly the same implications as Chebyshev polynomials or VFI.

Finally, Table 5 presents the welfare cost of the business cycle. Now, in comparison with the benchmark calibration, the welfare cost of the business cycle is positive and significant, slightly above 1.1 percent. This is not a surprise, since we have both a large risk aversion and productivity shocks with an average standard deviation three times as big as the observed one. All methods deliver numbers that are close.

Author's personal copy

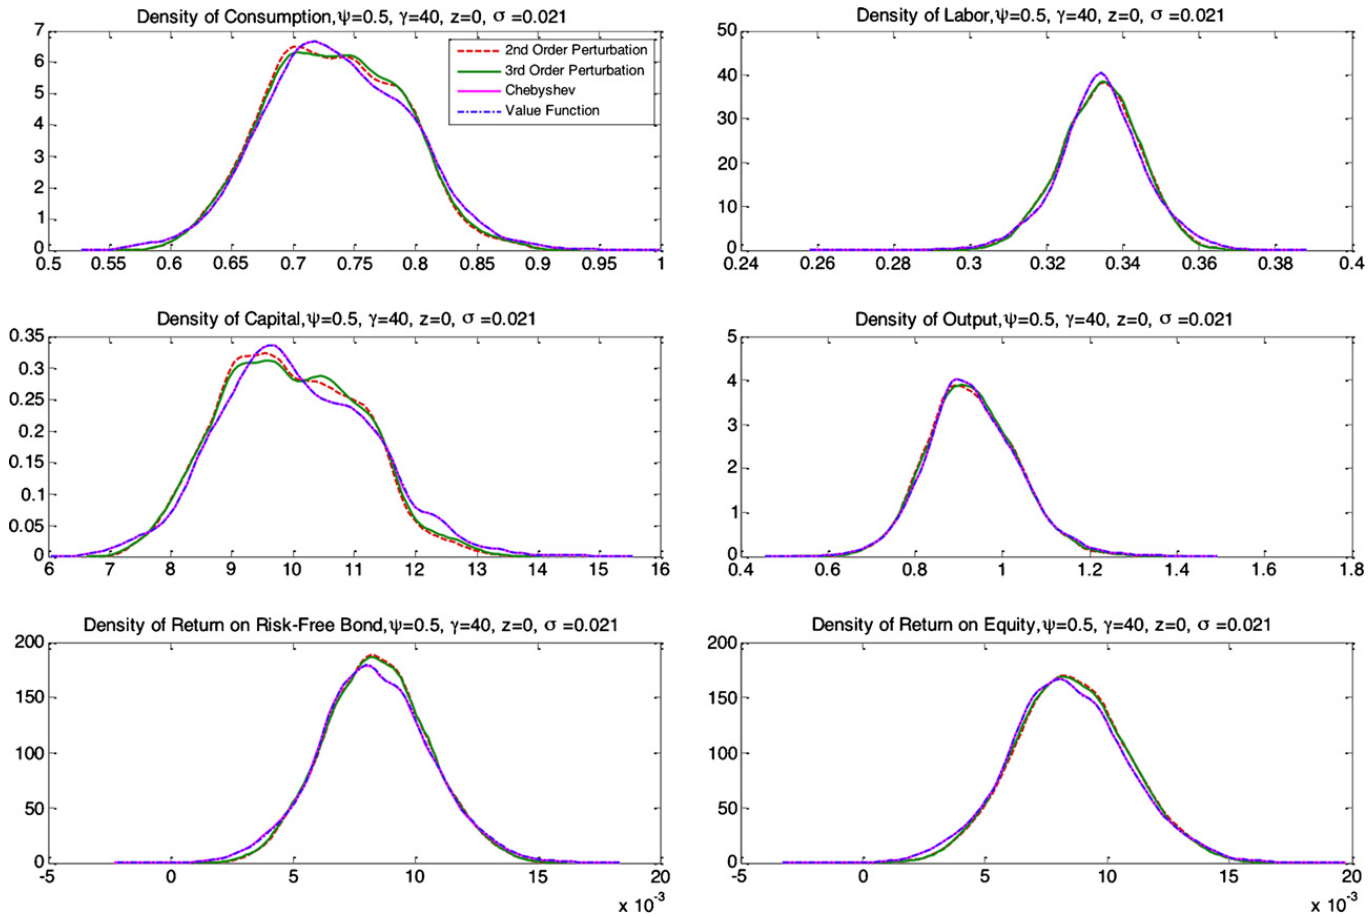

Fig. 4. Densities, extreme calibration.

Table 4
Business cycle statistics – benchmark calibration.

	c	y	i	R^f (%)	R^k (%)
	Mean				
Second-order perturbation	0.7338	0.9297	0.1950	0.8432	0.8562
Third-order perturbation	0.7344	0.9311	0.1955	0.8416	0.8529
Chebyshev polynomial	0.7359	0.9329	0.1970	0.8331	0.8402
Value function iteration	0.7359	0.9329	0.1970	0.8352	0.8403
	Variance (%)				
Second-order perturbation	0.2956	1.0575	0.2718	0.0004	0.0004
Third-order perturbation	0.3634	1.2178	0.3113	0.0004	0.0005
Chebyshev polynomial	0.3413	1.1523	0.3425	0.0005	0.0006
Value function iteration	0.3414	1.1528	0.3427	0.0005	0.0006

Table 5
Welfare costs of business cycle – extreme calibration.

2nd-order pert.	3rd-order pert.	Chebyshev	Value function
1.1278e−2	1.1278e−2	1.2855e−2	1.2838e−2

5.3. Euler equation errors

While the plots of the decision rules and the computation of densities and business cycle statistics that we presented in the previous subsection are highly informative, it is also important to evaluate the accuracy of each of the procedures. Euler equation errors, introduced by Judd (1992), have become a common tool for determining the quality of the solution method. The idea is to observe that, in our model, the intertemporal condition:

$$u_{c,t} = \beta\left(\mathbb{E}_t V_{t+1}^{1-\gamma}\right)^{\frac{1}{\theta}-1} \mathbb{E}_t\left(V_{t+1}^{\frac{(\gamma-1)(1-\theta)}{\theta}} u_{c,t+1} R(k_t, z_t, \sigma_t; z_{t+1}, \sigma_{t+1})\right) \tag{4}$$

Author's personal copy

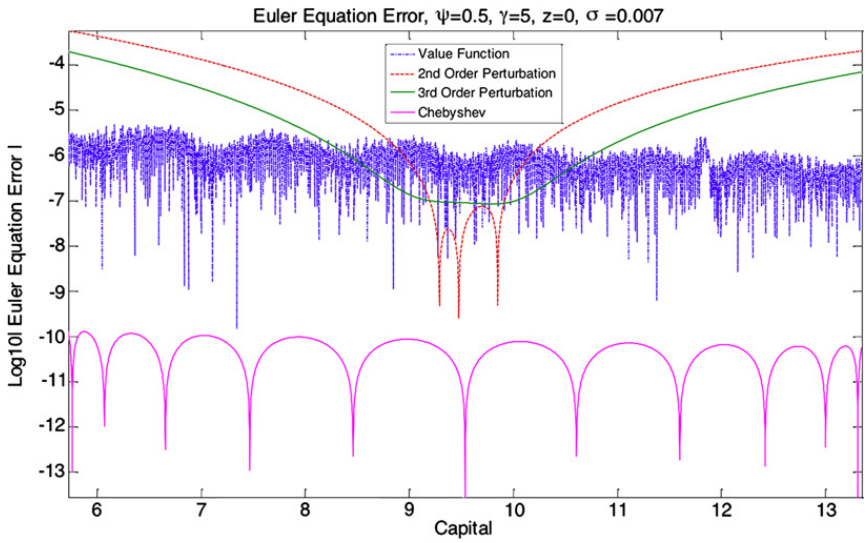

Fig. 5. Euler equation error, benchmark calibration.

where $R(k_t, z_t, \sigma_t; z_{t+1}, \sigma_{t+1}) = 1 + \zeta e^{z_{t+1}} k_{t+1}^{\zeta-1} l_{t+1}^{1-\zeta} - \delta$ is the gross return of capital given states k_t, z_t, σ_t, and realizations z_{t+1} and σ_{t+1} should hold exactly for any given k_t and z_t. However, since the solution methods we use are only approximations, there will be an error in (4) when we plug in the computed decision rules. This Euler equation error function $EE^i(k_t, z_t, \sigma_t)$ is defined, in consumption terms:

$$EE^i(k_t, z_t, \sigma_t) = 1 - \frac{\left[\dfrac{\beta\left(\mathbb{E}_t\left(V_{t+1}^i\right)^{1-\gamma}\right)^{\frac{1}{\theta}-1}\mathbb{E}_t\left(\left(V_{t+1}^i\right)^{\frac{(\gamma-1)(1-\theta)}{\theta}} u_{c,t+1}^i R(k_t, z_t, \sigma_t; z_{t+1}, \sigma_{t+1})\right)}{\frac{1-\gamma}{\theta}\upsilon\left(1-l_t^i\right)^{(1-\upsilon)\frac{1-\gamma}{\theta}}}\right]^{\frac{1}{\upsilon\frac{1-\gamma}{\theta}-1}}}{c_t^i}.$$

This function determines the (unit free) error in the Euler equation as a fraction of the consumption given the current states and solution method i. Following Judd and Guu (1997), we can interpret this error as the optimization error incurred by the use of the approximated decision rule and we report the absolute errors in base 10 logarithms to ease interpretation. Thus, a value of -3 means a \$1 mistake for each \$1000 spent, a value of -4 a \$1 mistake for each \$10,000 spent, and so on.

Fig. 5 displays a transversal cut of the errors for the benchmark calibration when $z = 0$ and $\sigma_t = \bar{\sigma}$. Other transversal cuts at different technology and volatility levels reveal similar patterns. The first lesson from Fig. 5 is that all methods deliver high accuracy. We know from Fig. 3 that capital is nearly always between 8.5 and 10.5. In that range, the ($\log 10$) Euler equation errors are at most -5, and most of the time they are even smaller. For instance, the second- and third-order perturbations have an Euler equation error of around -7 in the neighborhood of the deterministic steady state, VFI of around -6.5, and Chebyshev an impressive $-11/-13$. The second lesson from Fig. 5 is that, as expected, global methods (Chebyshev and VFI) perform very well in the whole range of capital values, while perturbations deteriorate as we move away from the steady state. For second-order perturbation, the Euler error in the steady state is almost four orders of magnitude smaller than on the boundaries. Third-order perturbation is around half an order of magnitude more accurate than second-order perturbation over the whole range of values (except in a small region close to the deterministic steady state).

There are two complementary ways to summarize the information from Euler equation error functions. First, we report the maximum error in our interval (capital between 60 percent and 140 percent of the steady state and the grids for productivity and volatility). The maximum Euler error is useful because it bounds the mistake owing to the approximation. The second procedure for summarizing Euler equation errors is to integrate the function with respect to the ergodic distribution of capital and productivity to find the average error. We can think of this exercise as a generalization of the Den Haan–Marcet test (Den Haan and Marcet, 1994). The top-left panel in Fig. 6 reports the maximum Euler error (darker bars) and the integral of the Euler error for the benchmark calibration. Both perturbations have a maximum Euler error of around -2.7, VFI of -3.1, and Chebyshev, an impressive -9.8. We read this result as indicating that all methods perform adequately. Both perturbations have roughly the same integral of the Euler error (around -5.3), VFI a slightly better -6.4, while Chebyshev polynomials do fantastically well at -10.4 (the average loss of welfare is \$1 for each \$500 billion). But even an approximation with an average error of \$1 for each \$200,000, such as the one implied by third-order perturbation, must suffice for most relevant applications.

We repeat our exercise for the extreme calibration. As we did when we computed the decision rules of the agents, we have changed the capital interval to $[3, 32]$. The top-right panel in Fig. 6 reports maximum Euler equation errors and their integrals. The maximum Euler equation error is large for perturbation methods while it is rather small using Chebyshev polynomials. However, given the very large range of capital used in the computation, this maximum Euler error provides a

Author's personal copy

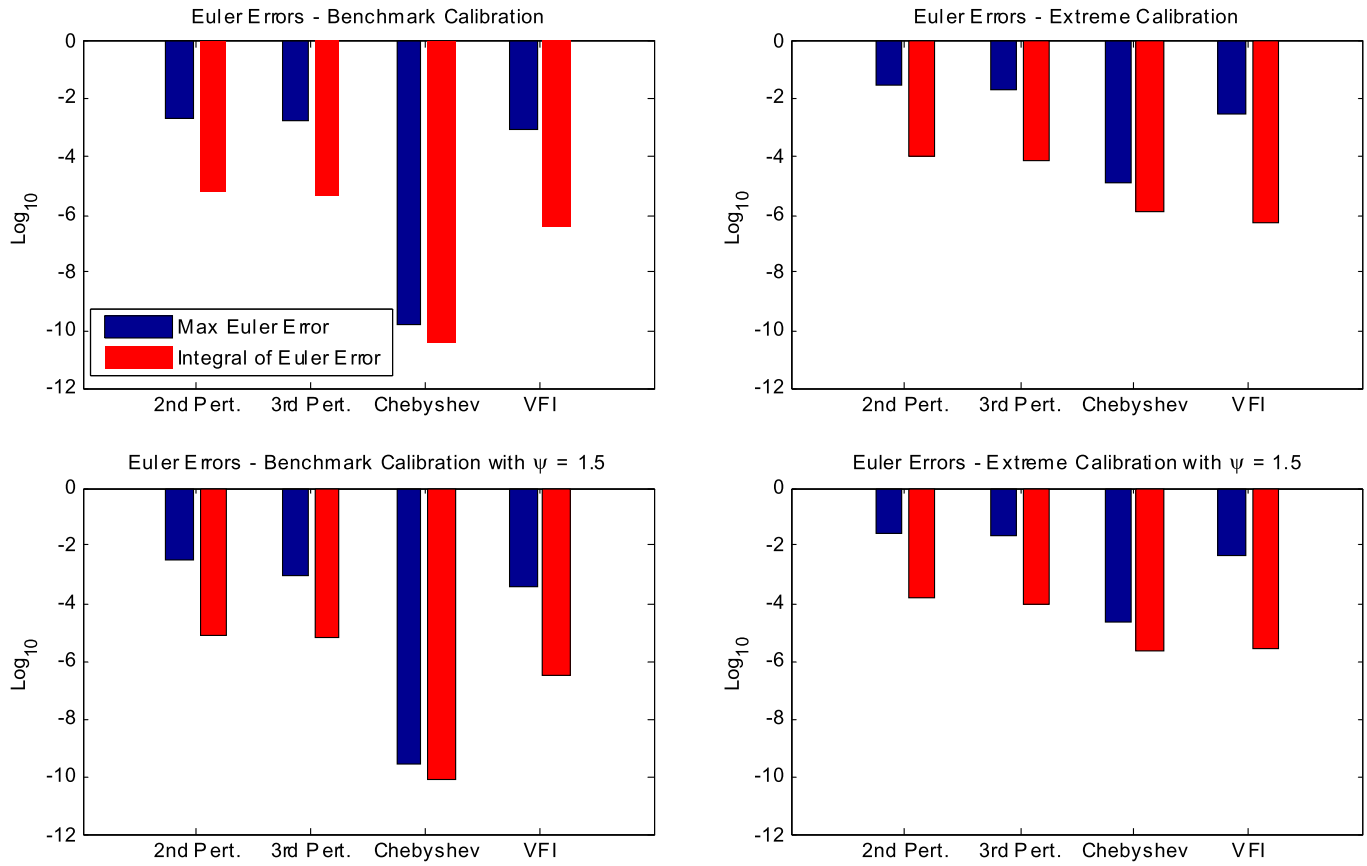

Fig. 6. Maximum Euler equation errors.

too negative view of accuracy. We find the integral of the Euler equation error to be more instructive. With a second-order perturbation, we have -4.02 and with a third-order perturbation we have -4.12. To evaluate this number, remember that we have extremely high risk aversion and large productivity shocks. Even in this challenging environment, perturbations deliver a high degree of accuracy. VFI does not display a big loss of precision compared to the benchmark case. On the other hand, Chebyshev polynomials deteriorate somewhat, but the accuracy it delivers it is still of \$1 out of each \$1 million spent.

5.4. Robustness: Changing the EIS and changing the perturbation point

In the results we reported above, we kept the EIS equal to 0.5, a conventional value in the literature, while we modified the risk aversion and the volatility of productivity shocks. However, since some researchers prefer higher values of the EIS (see, for instance, Bansal and Yaron, 2004, a paper that we have used to motivate our investigation), we also computed our model with $\psi = 1.5$. Basically our results were unchanged. To save on space, we concentrate only on the Euler equation errors (decision rules and simulation paths are available upon request). In the bottom-left panel in Fig. 6, we report the maxima of the Euler equation errors and their integrals with respect to the ergodic distribution. The relative size and values of the entries in this table are quite similar to the values reported for the benchmark calibration (except, partially, VFI that performs a bit better). The bottom-right panel in Fig. 6 repeats the same exercise for the extreme calibration. Again, the entries in the table are very close to the ones in the extreme calibration (and now, VFI does not perform better than when $\psi = 0.5$).

As a final robustness test, we computed the perturbations not around the deterministic steady state (as we did in the main text), but around a point close to the mode of the ergodic distribution of capital. This strategy, if perhaps difficult to implement because of the need to compute the mode of the ergodic distribution,[6] could deliver better accuracy because we approximate the value function and decision rules in a region where the model spends more time. As we suspected, we found only trivial improvements in terms of accuracy. Moreover, expanding at a point different from the deterministic

[6] For example, the algorithm of finding a perturbation around the steady state, simulated from it, finds a second perturbation around the model of the implied ergodic simulation, and so on until convergence, may not settle in any fixed point. In our exercise, we avoid this problem because we have the ergodic distribution implied by VFI. This is an unfair advantage for perturbations at the mode of the ergodic distribution but it makes our point below about the lack of improvement in accuracy even stronger.

Author's personal copy

steady state has the disadvantage that the theorems that ensure the convergence of the Taylor approximation might fail (see Theorem 6 in Jin and Judd, 2002).

5.5. Implementation and computing time

We briefly discuss implementation and computing time. For the benchmark calibration, second-order perturbation and third-order perturbation algorithms take only 0.02 second and 0.05 second, respectively, in a 3.3 GHz Intel PC with Windows 7 (the reference computer for all times below), and it is simple to implement: 664 lines of code in Fortran 95 for second order and 1133 lines of code for third order, plus in both cases, the analytical derivatives of the equilibrium conditions that Fortran 95 borrows from a code written in `Mathematica 6`.[7] The code that generates the analytic derivatives has between 150 to 210 lines, although `Mathematica` is much less verbose. While the number of lines doubles in the third order, the complexity in terms of coding does not increase much: the extra lines are mainly from declaring external functions and reading and assigning values to the perturbation coefficients. An interesting observation is that we only need to take the analytic derivatives once, since they are expressed in terms of parameters and not in terms of parameter values. This allows Fortran to evaluate the analytic derivatives extremely fast for new combinations of parameter values. This advantage of perturbation is particularly relevant when we need to solve the model repeatedly for many different parameter values, for example, when we are estimating the model. For completeness, the second-order perturbation was also run in `Dynare` (although we had to use version 4.0, which computes analytic derivatives, instead of previous versions, which use numerical derivatives that are not accurate enough for perturbation). This run was a double-check of the code and a test of the feasibility of using off-the-shelf software to solve DSGE models with recursive preferences.

The projection algorithm takes around 300 seconds, but it requires a good initial guess for the solution of the system of equations. Finding the initial guess for some combination of parameter values proved to be challenging. The code is 652 lines of Fortran 95. Finally, the VFI code is 707 lines of Fortran 95, but it takes about ten hours to run.

6. Conclusions

In this paper, we have compared different solution methods for DSGE models with recursive preferences and SV. We evaluated the different algorithms based on accuracy, speed, and programming burden. We learned that all of the most promising methods (perturbation, projection, and VFI) do a fair job in terms of accuracy. We were surprised by how well simple second-order and third-order perturbations perform even for fairly non-linear problems. We were impressed by how accurate Chebyshev polynomials can be. However, their computational cost was higher and we are concerned about the curse of dimensionality. In any case, it seems clear to us that, when accuracy is the key consideration, Chebyshev polynomials are the way to go. Finally, we were disappointed by VFI since even with 125,000 points in the grid, it only did marginally better than perturbation and it performed much worse than Chebyshev polynomials in our benchmark calibration. This suggests that unless there are compelling reasons such as non-differentiabilities or non-convexities in the model, we better avoid VFI.

A theme we have not developed in this paper is the possibility of interplay among different solution methods. For instance, we can compute extremely easily a second-order approximation to the value function and use it as an initial guess for VFI. This second-order approximation is such a good guess that VFI will converge in few iterations. We verified this idea in non-reported experiments, where VFI took one-tenth of the time to converge once we used the second-order approximation to the value function as the initial guess. This approach may even work when the true value function is not differentiable at some points or has jumps, since the only goal of perturbation is to provide a good starting point, not a theoretically sound approximation. This algorithm may be particularly useful in problems with many state variables. More research on this type of hybrid method is a natural extension of our work.

We close the paper by pointing out that recursive preferences are only one example of a large class of non-standard preferences that have received much attention by theorists and applied researchers over the last several years (see Backus et al., 2004). Having fast and reliable solution methods for this class of new preferences will help researchers to sort out which of these preferences deserve further attention and to derive empirical implications. Thus, this paper is a first step in the task of learning how to compute DSGE models with non-standard preferences.

Appendix A

In this appendix, we present the steady state of the model and the alternative perturbation approach, the value function perturbation (VFP).

A.1. Steady state of the model

To solve the system:

[7] We use lines of code as a proxy for the complexity of implementation. We do not count comment lines.

Author's personal copy

$$V_{ss} = c_{ss}^{\upsilon}(1-l_{ss})^{1-\upsilon}$$
$$\left(\zeta k_{ss}^{\zeta-1} l_{ss}^{1-\zeta} + 1 - \delta\right) = 1/\beta$$
$$\frac{1-\nu}{\nu}\frac{c_{ss}}{(1-l_{ss})} = (1-\zeta)k_{ss}^{\zeta} l_{ss}^{-\zeta}$$
$$m_{ss} R_{ss}^{f} = 1/\beta$$
$$c_{ss} + i_{ss} = k_{ss}^{\zeta} l_{ss}^{1-\zeta}$$
$$i_{ss} = \delta k_{ss}$$

note first that:

$$\frac{k_{ss}}{l_{ss}} = \left(\frac{1}{\zeta}\left(\frac{1}{\beta} - 1 + \delta\right)\right)^{\frac{1}{\zeta-1}} = \Omega.$$

Now, from the leisure-consumption condition:

$$\frac{c_{ss}}{1-l_{ss}} = \frac{\upsilon}{1-\upsilon}(1-\zeta)\Omega^{\zeta} = \Phi \quad \Rightarrow \quad c_{ss} = \Phi(1-l_{ss}).$$

Then:

$$c_{ss} + \delta k_{ss} = k_{ss}^{\zeta} l_{ss}^{1-\zeta} = \Omega^{\zeta} l_{ss} \quad \Rightarrow \quad c_{ss} = \left(\Omega^{\zeta} - \delta\Omega\right) l_{ss}$$

and:

$$\Phi(1-l_{ss}) = \left(\Omega^{\zeta} - \delta\Omega\right) l_{ss} \quad \Rightarrow$$
$$l_{ss} = \frac{\Phi}{\Omega^{\zeta} - \delta\Omega + \Phi}$$
$$k_{ss} = \frac{\Phi\Omega}{\Omega^{\zeta} - \delta\Omega + \Phi}$$

from which we can find V_{ss} and i_{ss}.

A.2. Value function perturbation (VFP)

We mentioned in the main text that instead of perturbing the equilibrium conditions of the model, we could directly perturb the value function in what we called value function perturbation (VFP). To undertake the VFP, we write the value function as:

$$V(k_t, z_t, \sigma_t; \chi) = \max_{c_t, l_t}\left[(1-\beta)\left(c_t^{\upsilon}(1-l_t)^{1-\upsilon}\right)^{\frac{1-\gamma}{\theta}} + \beta\mathbb{E}_t V^{1-\gamma}(k_{t+1}, z_{t+1}, \sigma_{t+1}; \chi)^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}.$$

To find a second-order approximation to the value function, we take derivatives of the value function with respect to controls (c_t, l_t), states (k_t, z_t, σ_t), and the perturbation parameter χ. We collect these 6 equations, together with the resource constraint, the value function itself, and the exogenous processes in a system:

$$\widetilde{F}(k_t, z_t, \chi) = \mathbf{0}$$

where the hat over F emphasizes that now we are dealing with a slightly different set of equations than the F in the main text.

After solving for the steady state of this system, we take derivatives of the function $\widetilde{F}$ with respect to k_t, z_t, σ_t, and χ:

$$\widetilde{F}_i(k_{ss}, 0, \sigma_{ss}; 0) = \mathbf{0} \quad \text{for } i = \{1, 2, 3, 4\}$$

and we solve for the unknown coefficients. This solution will give us a second-order approximation of the value function but only a first-order approximation of the decision rules. By repeating these steps n times, we can obtain the $(n+1)$-order approximation of the value function and the n-order approximation of the decision rules. It is straightforward to check that the coefficients obtained by ECP and VFP are the same. Thus, the choice of one approach or the other should be dictated by expediency.

Author's personal copy

References

Anderson, E.W., Hansen, L.P., McGrattan, E.R., Sargent, T.J., 1996. On the mechanics of forming and estimating dynamic linear economies. In: Amman, H.M., et al. (Eds.), Handbook of Computational Economics. Elsevier.

Aruoba, S.B., Fernández-Villaverde, J., Rubio-Ramírez, J., 2006. Comparing solution methods for dynamic equilibrium economies. Journal of Economic Dynamics and Control 30, 2477–2508.

Backus, D.K., Routledge, B.R., Zin, S.E., 2004. Exotic preferences for macroeconomists. NBER Macroeconomics Annual 2004, 319–390.

Backus, D.K., Routledge, B.R., Zin, S.E., 2007. Asset pricing in business cycle analysis. Mimeo. New York University.

Bansal, R., Yaron, A., 2004. Risks for the long run: A potential resolution of asset pricing puzzles. Journal of Finance 59, 1481–1509.

Benigno, P., Woodford, M., 2006. Linear-quadratic approximation of optimal policy problems. Mimeo. Columbia University.

van Binsbergen, J.H., Fernández-Villaverde, J., Koijen, R.S.J., Rubio-Ramírez, J., 2009. Likelihood estimation of DSGE models with Epstein–Zin preferences. Mimeo. University of Pennsylvania.

Cho, J.-O., Cooley, T.C., 2000. Business cycle uncertainty and economic welfare. Mimeo. New York University.

Chow, C.S., Tsitsiklis, J.N., 1991. An optimal one-way multigrid algorithm for discrete-time stochastic control. IEEE Transaction on Automatic Control 36, 898–914.

Christiano, L., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1–45.

Den Haan, W.J., Marcet, A., 1994. Accuracy in simulations. Review of Economic Studies 61, 3–17.

Epstein, L., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica 57, 937–969.

Epstein, L., Zin, S.E., 1991. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. Journal of Political Economy 99, 263–286.

Fernández-Villaverde, J., Rubio-Ramirez, J., 2006. Solving DSGE models with perturbation methods and a change of variables. Journal of Economic Dynamics and Control 30, 2509–2531.

Fernández-Villaverde, J., Rubio-Ramirez, J., 2010. Macroeconomics and volatility: Data, models, and estimation. NBER Working Paper 16618.

Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramírez, J., 2010. Fortune or virtue: Time-variant volatilities versus parameter drifting in U.S. data. NBER Working Paper 15928.

Fornberg, B., 1998. A Practical Guide to Pseudospectral Methods. Cambridge University Press, Cambridge.

Hansen, L.P., Sargent, T.J., 1995. Discounted linear exponential quadratic Gaussian control. IEEE Transactions on Automatic Control 40, 968–971.

Hansen, L.P., Heaton, J., Lee, J., Roussanov, N., 2007. Intertemporal substitution and risk aversion. In: Heckman, James J., Leamer, Edward E. (Eds.), Handbook of Econometrics, vol. 6, pp. 3967–4056.

Jin, H., Judd, K.L., 2002. Perturbation methods for general dynamic stochastic models. Mimeo. Hoover Institution.

Judd, K.L., 1992. Projection methods for solving aggregate growth models. Journal of Economic Theory 58, 410–452.

Judd, K.L., 1998. Numerical Methods Economics. MIT Press.

Judd, K.L., Guu, S.M., 1992. Perturbation solution methods for economic growth models. In: Varian, H. (Ed.), Economic and Financial Modelling with Mathematica. Springer-Verlag.

Judd, K.L., Guu, S.M., 1997. Asymptotic methods for aggregate growth models. Journal of Economic Dynamics and Control 21, 1025–1042.

Kreps, D.M., Porteus, E.L., 1978. Temporal resolution of uncertainty and dynamic choice theory. Econometrica 46, 185–200.

Levine, P., Pearlman, J., Pierse, R., 2007. Linear-quadratic approximation, external habit, and targeting rules. Mimeo. University of Surrey.

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.

Schmitt-Grohé, S., Uribe, M., 2005. Optimal fiscal and monetary policy in a medium scale macroeconomic model. NBER Macroeconomic Annual 2005, 382–425.

Smets, F., Wouters, R., 2007. Shocks and frictions in US business cycles: A Bayesian DSGE approach. American Economic Review 97, 586–606.

Tallarini, T.D., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45, 507–532.

Tauchen, G., 1986. Finite state Markov–Chain approximations to univariate and vector autoregressions. Economics Letters 20, 177–181.

Weil, P., 1990. Nonexpected utility in macroeconomics. Quarterly Journal of Economics 105, 29–42.

Woodford, M.D., 2003. Interest and Prices. Princeton University Press.