ISSN 1518-3548

BANCO CENTRAL DO BRASIL

Working Paper Series 281

A Note on Particle Filters Applied to DSGE Models

Angelo Marsiglia Fasolo

June, 2012

ISSN 1518-3548
CGC 00.038.166/0001-05

Working Paper Series	Brasília	n. 281	Jun.	2012	p. 1-36

Working Paper Series

Edited by Research Department (Depep) – E-mail: workingpaper@bcb.gov.br

Editor: Benjamin Miranda Tabak – E-mail: benjamin.tabak@bcb.gov.br
Editorial Assistant: Jane Sofia Moita – E-mail: jane.sofia@bcb.gov.br
Head of Research Department: Adriana Soares Sales – E-mail: adriana.sales@bcb.gov.br

The Banco Central do Brasil Working Papers are all evaluated in double blind referee process.

Authorized by Carlos Hamilton Vasconcelos Araújo, Deputy Governor for Economic Policy.

General Control of Publications

Banco Central do Brasil

Secre/Comun/Cogiv

SBS – Quadra 3 – Bloco B – Edifício-Sede – 1º andar

Caixa Postal 8.670

70074-900 Brasília – DF – Brazil

Phones: +55 (61) 3414-3710 and 3414-3565

Fax: +55 (61) 3414-3626

E-mail: editor@bcb.gov.br

Consumer Complaints and Public Enquiries Center

Banco Central do Brasil

Secre/Surel/Diate

SBS – Quadra 3 – Bloco B – Edifício-Sede – 2º subsolo

70074-900 Brasília – DF – Brazil

Fax: +55 (61) 3414-2553

Internet: <http://www.bcb.gov.br/?english>

A Note on Particle Filters Applied to DSGE Models

Angelo Marsiglia Fasolo*

Abstract

This paper compares the properties of two particle filters — the Bootstrap Filter and the Auxiliary Particle Filter — applied to the computation of the likelihood of artificial data simulated from a basic DSGE model with nominal and real rigidities. Particle filters are compared in terms of speed, quality of the approximation of the probability density function of data and tracking of state variables. Results show that there is a case for the use of the Auxiliary Particle Filter only when the researcher uses a large number of observable variables and the number of particles used to characterize the likelihood is relatively low. Simulations also show that the largest gains in tracking state variables in the model are found when the number of particles is between 20,000 and 30,000, suggesting a boundary for this number.

JEL Codes: C63, C68, E30, E37

Keywords: Particle filters; Bootstrap filter; Auxiliary particle filter; Non-linear methods; DSGE models

*Research Department, Banco Central do Brasil. The author is thankful for the comments of Juan Rubio-Ramírez and Pablo Guerrón-Quintana on the implementation of the particle filter. The author would also like to thank comments of Hernán D. Seoane and the participants of seminar at UFRGS and Central Bank of Brazil. Remaining errors, however, are entirely due to author's fault. E-mail: angelo.fasolo@bcb.gov.br

1 Introduction

Models in state-space representation constitute a benchmark in today's economic analysis. The widespread use of linearized representation of equilibrium conditions made the Kalman Filter an important tool for simulation of non-observed variables, helping on the estimation of economic models. Particularly in the analysis of Dynamic Stochastic General Equilibrium (DSGE) Models, the seminal contribution of Smets and Wouters (2003)[28] quickly became a standard methodology for estimating medium and large scale DSGE models, with several applications in current literature. The authors rely on the Kalman Filter to obtain the likelihood of a model characterizing the European economy and estimate it using Bayesian techniques. The Kalman Filter, however, assumes a structure sometimes very restrictive if the researcher is interested in non-linear phenomena usually observed in data[1]. Additionally, Fernández-Villaverde, Rubio-Ramírez and Santos (2006)[15] show that the errors in the approximation of equilibrium conditions of a model accumulate over the sample. As a consequence, second-order errors of approximated policy functions still generate first-order changes in the approximation of the likelihood of a model[2]. In order to overcome these problems, Fernández-Villaverde and Rubio-Ramírez (2005, 2007)[13] [14] propose the use of a particle filter to compute the likelihood of DSGE models, generalizing the methodology described in Smets and Wouters (2003)[28]. This paper characterizes two particle filters applied on a standard DSGE model in terms of speed, the quality of the likelihood approximation and the ability to track the model's state variables. For all those features, the number of particles characterizing the distribution of state variables plays a significant role.

Particle filters, also called Sequential Importance Sampling methods or Sequential Monte Carlo methods, constitute an important set of procedures used in a large variety of applications in different fields of science. In economics, however, only recently the Bootstrap Filter – a particular implementation of particle filters –, proposed by Gordon, Salmond and Smith (1993)[18], got some attention as a statistical tool to evaluate the likelihood function in complex, non-linear models. As an example, Kim, Shephard and Chib (1998) and Pitt and Shephard (1999)[25] study simple models with stochastic volatility. The large number of state variables in standard models, combined with the large computational burden associated with frequent calls to the filter in simulation procedures, resulted in prohibitive costs for a practical implementation. With the increased use of parallel distribution of tasks in personal computers with multiple CPU cores and GPU processing, the time constraints on procedures based on particle filtering are vastly reduced.

The consistency of the particle filter is based on the number of particles used to approximate

[1] Among others, monetary and fiscal policy shifts of regime, the "fat-tails" usually verified in economic shocks, the influence of risk aversion and precautionary savings on aggregates like consumption and investment are all phenomena not well characterized in a linear framework. The combination of a linear state-space representation with Gaussian disturbances – the basic framework underlying the Kalman Filter – is not appropriate to handle these features.

[2] Alves (2011)[1] simulates a very long artificial dataset from a non-linear solution of the New-Keynesian model and estimates the structural parameters using log-linear approximations of the equilibrium conditions. He finds significant bias in the estimation of the labor supply elasticity and the price rigidity (Calvo) parameters.

the probability density function of state variables in each point in time, at the expense of additional time spent in computations. The trade-off between accuracy of the filter and the time necessary to execute the procedure can be summarized by results in figure 1. In order to build the results presented in the figure, the model in Christiano, Eichenbaum and Evans (2005)[9] was simulated to generate an artificial data with 100 observations. Using the artificial data and the true parameter set generating data, the Bootstrap Filter simulated the log-likelihood of the model 1000 times for different sizes of the particle swarm. The only source of uncertainty in the simulations is the random draw generating states.

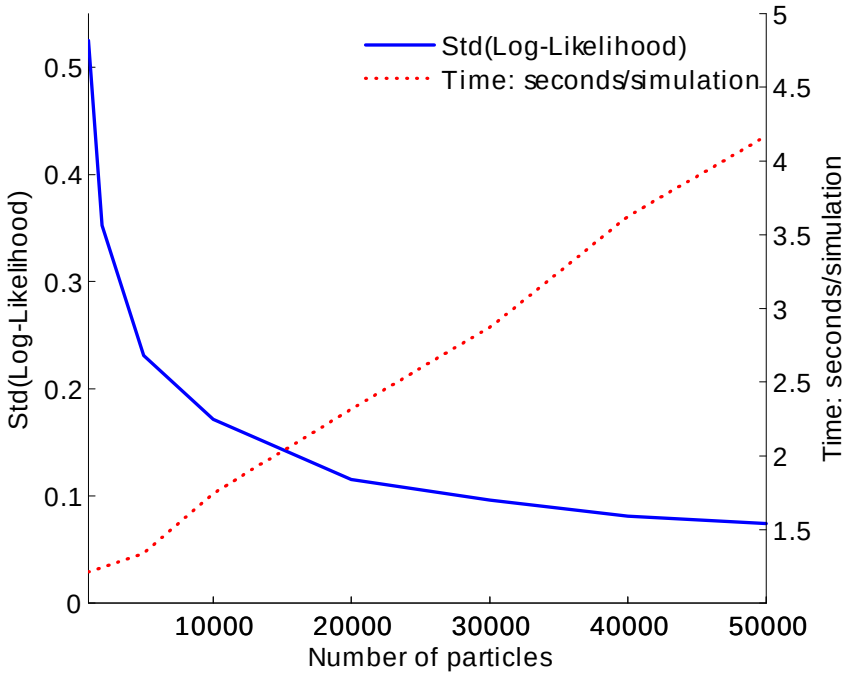

Figure 1: Bootstrap Filter: Standard-Deviation of Log-Likelihood and Time per Simulation

In figure 1, the horizontal axis shows the number of particles used to generate the simulations. The vertical axis on the left shows the standard deviation of the 1000 draws of the log-likelihood, while the axis on the right shows the average time per simulation measured in seconds. The figure shows the importance of selecting an appropriate number of particles on the set up of particle filters: simulated values of the log-likelihood become more precise as the number of particles increase. However, as the gains in precision declines with the increase in the number of particles, the average time per simulation, and the computational burden associated, increases over an almost linear trend.

It is interesting, though, that there is not a clear setup on how to determine the appropriate number of particles to use for likelihood inference. Indeed, there is a large discrepancy in

literature regarding the relation between the size of the model – measured in terms of number of state variables – and the number of particles used for simulation purposes. There are very few papers specifically addressing this problem. One is presented in Fernández-Villaverde and Rubio-Ramírez (2004)[12], where the authors simulate 50 times the log-likelihood of a model under alternative calibrations. Using 40,000 particles, the standard-deviation of the simulated log-likelihood is less than half-percentage point of the mean value of the log-likelihood. Andreasen (2011)[4] proposes a new distribution for the importance sampling based on the Central Difference Kalman Filter and tests his version of the particle filter against the basic filter described in Fernández-Villaverde and Rubio-Ramírez (2004)[12] and two alternative versions. The author shows that his filter outperforms the basic particle filter in terms of standard-deviation of the simulated log-likelihood for models with a small number of shocks and a low number of particles. In a New Keynesian model with 3 and 4 shocks and with 50,000 particles, the standard-deviation of simulations become very similar across filters.

In empirical applications, Fernández-Villaverde and Rubio-Ramírez (2005)[13] simulate the likelihood function of Neoclassical Growth Model using 60,000 particles. Fernández-Villaverde and Rubio-Ramírez (2007)[14] estimate an extended version of the Neoclassical Growth Model with investment-specific technological change, preference shocks and stochastic volatility in shocks. The extended version contains eight state variables against only two from the basic model. The estimation used 80,000 particles in the setup of the particle filter. An and Schorfheide (2007)[3] and An (2008)[2] estimate the basic New-Keynesian model with quadratic adjustment costs for prices using US data. The Bootstrap Filter is set with 40,000 particles. Lombardi and Sgherri (2007)[24] set a simple DSGE model without capital and use the particle filter to estimate the model and track the path of the natural interest rate. The filter is set with only 10,000 particles. Fernández-Villaverde, Guerrón-Quintana and Rubio-Ramírez (2010)[11] estimate a fully-fledged DSGE model with sticky prices and wages, plus parameter drifts and stochastic volatility to shocks, for a total of 19 state variables. For this large model, the likelihood is computed using only 10,000 particles, mainly due to speed concerns when running the code. Flury and Shephard (2011)[16] apply the particle filter for the estimation of several types of economic models. They set the particle filter with 60,000 particles when simulating the likelihood of a very simple DSGE model without nominal frictions.

This paper addresses the problem of selecting an appropriate number of particles while comparing two specific particle filters applied to DSGE models: the Bootstrap Filter, proposed by Gordon, Salmond and Smith (1993)[18], and the extension known as the Auxiliary Particle Filter, proposed by Pitt and Shephard (1999)[25]. The problem of selecting the size of the particle swarm is relevant not only in terms of properly characterizing the density of state variables, but also due to the problem of sample depletion in common settings adopted in the estimation of DSGE models. As will be shown, the problem of sample depletion manifests not only as a consequence of an insufficient number of particles, but also when independent measurement errors of observable variables – a priori unknown to the researcher – converges to zero. Most exercises

are based on the computation of the number of effective particles proposed in Liu and Chen (1995)[23], combined with the time necessary to compute the likelihood function. Additionally, the ability of each filter to correctly track the path of state variables, emphasized by Boers (1999)[7], is evaluated in each simulation. Robustness is checked changing the model for number of observable and state variables. The Christiano, Eichenbaum and Evans (2005)[9] model was selected because it is a well-known benchmark in terms of DSGE models. It has a relatively large number of state variables and provides enough flexibility to alter the number of state and observable variables with small changes in the baseline calibration. This flexibility is extensively explored during robustness tests.

Note that, when discussing about an "appropriate number of particles", this paper is not presenting a procedure to optimaly determine the size of the particle swarm. Instead, the approach here consists in evaluating the ability of each filter in correctly tracking state variables of the model. As shown in Boers (1999)[7], statistics related to the prediction error of state variables stabilize after a given particle swarm size. Given the linear expansion of computational time as the swarm gets larger, it is natural to not spend additional resources when gains become limited. As a consequence, the appropriate number of particles can also be interpreted as the minimum size of the particle swarm required to properly work with the particle filter.

In terms of results, this paper shows that the Auxiliary Particle Filter provides better results over the Bootstrap Filter when the size of the particle swarm is small or the number of observable variables is relatively large. The gains of the Auxiliary Particle Filter are mostly concentrated at the tails of the distribution of the probability density function. However, as the size of the swarm increases, the signal-to-noise ratio of the model is reduced or the number of observable variables is reduced, there is not a clear advantage in using the Auxiliary Particle Filter, as the computational burden offsets gains in accuracy of the filter. As a basic setup of both filters, results suggest that a swarm between 20,000 and 30,000 particles has an adequate size to properly track the state variables of the model. This result is not conditional on the number of state variables or the size of measurement errors characterizing the model. The paper does not deal with complex proposals of particle filters, as in Andreasen (2011)[4], but instead shows that, based on the quality of the approximation of state variables, a large number of particles might not be necessary to properly characterize the density of states. Results also stress the role of GPU computing to calculate the likelihood of the model: in the model with 50,000 particles and three shocks, Andreasen (2011)[4] spends around 15 seconds in the basic filter for each evaluation of the likelihood for a sample with 200 observations; figure 1 shows that the average evaluation of the likelihood of a sample of 100 observations takes less than 4.5 seconds.

The paper is organized as follows. Section 2 provides a brief description of both the Bootstrap Filter and the Auxiliary Particle Filter. Section 3 presents the theoretical model, calibration and solution method. Section 4 presents the main results while section 5 concludes.

2 Two Particle Filters

In this section, two versions of the Particle Filter are presented: the Bootstrap Filter, which is a basic filter using a specific function for the importance sampling procedure; and the Auxiliary Particle Filter, an extension that includes more information when approximating the probability density function at the expense of additional computational time. The Bootstrap Filter has been used in the literature applied to DSGE models, notably at Fernández-Villaverde and Rubio-Ramírez (2005)[13], Fernández-Villaverde and Rubio-Ramírez (2007)[14] and Fernández-Villaverde, Guerrón-Quintana and Rubio-Ramírez (2010)[11]. The Auxiliary Particle Filter is an extension proposed by Pitt and Shephard (1999)[25] designed to incorporate information about the current period when filtering information. Most of its applications are based on small models due to the additional computational time.

Consider a model where the researcher has access to a vector of n observable variables[3] in each point in time $Y^T = \{Y_t\}_{t=1}^T \in \mathbb{R}^{nT}$ with the elements of Y^T given by $y^T = \{y_t\}_{t=1}^T$. There is a set of state variables, S_t, and parameters $\gamma \in \Upsilon$ characterizing the model which are not necessarily observed. The researcher would like to obtain a likelihood function given Y^T and a proposed parameter set γ :

$$L\left(Y^T;\gamma\right) = p\left(Y^T;\gamma\right)$$

The solution of the dynamic system is composed by two general functions g and f describing the evolution of state variables and relating observable and state variables:

$$Y_t = g\left(S_t, V_t;\gamma\right) \tag{1}$$

$$S_t = f\left(S_{t-1}, W_t;\gamma\right) \tag{2}$$

In the representation, $\{V_t\}$ and $\{W_t\}$ are sequences of exogenous, not necessarily independent variables. As a set of sequences, however, $\{V_t\}$ and $\{W_t\}$ are independent of each other. More specifically, $\{V_t\}$ is interpreted as noise to the set of observations, while $\{W_t\}$ is a set of exogenous shocks included in the characterization of the model. Note that, in principle, there are no assumptions regarding the functional forms g and f and the distribution of $\{V_t\}$ and $\{W_t\}$: it is only necessary that the probability density functions $p\left(W_t;\gamma\right)$ and $p\left(V_t;\gamma\right)$ are known and can be evaluated by the researcher. In terms of the size of the sequences, in order to avoid stochastic singularity in the model, it is necessary that $\dim\left(W_t\right) + \dim\left(V_t\right) \geq \dim\left(Y_t\right)$. Furthermore, to clarify the exposition, assume that $\{W_t\}$ can be partitioned in two sequences $\{W_{1,t}\}$ and $\{W_{2,t}\}$ such that $W_t = \left(W_{1,t}, W_{2,t}\right)$ and $\dim\left(W_{2,t}\right) + \dim\left(V_t\right) = \dim\left(Y_t\right)$. Also, to make notation clear, define $W_i^T = \{W_{i,t}\}_{t=1}^T$, for $i = 1, 2$, $V^T = \{V_t\}_{t=1}^T$ and $S^T = \{S_t\}_{t=1}^T$.

In order compute the likelihood, first note that, given the notation above, $L\left(Y^T;\gamma\right)$ can be

[3] For the sake of simplicity, notation in this section follows closely Fernández-Villaverde and Rubio-Ramírez (2005)[13] and Fernández-Villaverde and Rubio-Ramírez (2007)[14]. The reader is strongly encouraged to refer to Fernández-Villaverde and Rubio-Ramírez (2007)[14] for a more technical exposition of the problem.

factored as:

$$\begin{aligned} p\left(y^T;\gamma\right) &= \prod_{t=1}^{T} p\left(y_t|y_{t-1};\gamma\right) \\ &= \prod_{t=1}^{T} \iint p\left(y_t|W_1^t, S_0, y^{t-1};\gamma\right) p\left(W_1^t, S_0|y^{t-1};\gamma\right) dW_1^t dS_0 \end{aligned}$$

The first step simply defines the likelihood function as the product of densities evaluated in each period of time t, conditional on the information available in the previous period. The integrals in the second step describes the density in each point in time: it decomposes past information in terms of the sequence of exogenous variables and the initial state of the model, jointly integrated over time. Assume, for a moment, that g and f are linear functions of the state S_t, and $\{V_t\}$ and $\{W_t\}$ are both Normally distributed. In this specific case, the Kalman filter and its prediction error decomposition can be easily applied to obtain exact expressions for both integrals in terms of Riccati equations – the value of the likelihood $L\left(Y^T;\gamma\right)$ is exact, not an approximation[4]. In the general case proposed here, where there is not a closed form solution available, alternative methods are employed to obtain an approximation of $L\left(Y^T;\gamma\right)$.

The Bootstrap Filter allows a proper evaluation of the integrals, given that a sufficient number of simulations are performed. The key step here is the assumption about the knowledge of the distribution of $p\left(W_t;\gamma\right)$: given a swarm of particles $\left\{w_1^{t-1,i}, s_0^{t-1,i}\right\}_{i=1}^{N}$ of size N from the distribution in $t-1$ of $p\left(W_1^{t-1}, S_0|y^{t-1};\gamma\right)$, the distribution $p\left(W_{1,t};\gamma\right)$ is used to generate a set of draws characterizing the proposed distribution of $p\left(W_1^{t|t-1}, S_0|y^{t-1};\gamma\right)$, since $p\left(W_1^{t|t-1}, S_0|y^{t-1};\gamma\right) = p\left(W_{1,t};\gamma\right) p\left(W_1^{t-1}, S_0|y^{t-1};\gamma\right)$. Moving from a proposal to the actual distribution $p\left(W_1^t, S_0|y^{t-1};\gamma\right)$ is an application of Importance Sampling methods (Geweke, 1989[17]), using $p\left(W_1^{t|t-1}, S_0|y^{t-1};\gamma\right)$ as the so-called importance sampling distribution. Define the importance weights derived from the importance sampling distribution as the probability that a particle $\left(w_1^{t|t-1,i}, s_0^{t|t-1,i}\right)$ is chosen from the distribution $p\left(W_1^{t|t-1}, S_0|y^{t-1};\gamma\right)$:

$$q_t^i = \frac{p\left(y_t|w_1^{t|t-1,i}, s_0^{t|t-1,i}, y^{t-1};\gamma\right)}{\sum_{i=1}^{N} p\left(y_t|w_1^{t|t-1,i}, s_0^{t|t-1,i}, y^{t-1};\gamma\right)}$$

The final distribution $p\left(W_1^t, S_0|y^{t-1};\gamma\right)$ is a result of a resample procedure with replacement of a swarm of size N of $\left\{w_1^{t,i}, s_0^{t,i}\right\}_{i=1}^{N}$ being chosen with probability q_t^i from the proposal distribution in $p\left(W_1^{t|t-1}, S_0|y^{t-1};\gamma\right)$. The algorithm below summarizes the procedure of the Bootstrap filter:

[4] See Hamilton (1989)[19] for a detailed description of the Kalman filter and the prediction error decomposition used to calculate $L\left(Y^T;\gamma\right)$.

Algorithm 1: Bootstrap Filter

1) Initialization: $t = 0$. For $i = 1, ..., N$, initialize $p\left(W_1^{t-1}, S_0 | y^{t-1}; \gamma\right)$ and set $t = 1$.

2) Prediction: Generate N proposals of $\left\{w_1^{t-1,i}, s_0^{t-1,i}\right\}_{i=1}^{N}$ from the distribution $p\left(W_1^{t|t-1}, S_0 | y^{t-1}; \gamma\right) = p\left(W_{1,t}; \gamma\right) p\left(W_1^{t-1}, S_0 | y^{t-1}; \gamma\right)$.

3) Filtering: Evaluate importance weights q_t^i.

4) Sampling: Resample with replacement N draws from $\left\{w_1^{t|t-1,i}, s_0^{t|t-1,i}\right\}_{i=1}^{N}$ based on probabilities q_t^i. Call the draws $\left\{w_1^{t,i}, s_0^{t,i}\right\}_{i=1}^{N}$. If $t < T$, set $t \rightsquigarrow t + 1$ and go to step 2. Stop otherwise.

The use of the model to simulate next period's states, as defined in the prediction step, is the main characteristic of the Bootstrap Filter described in Gordon, Salmond and Smith (1993)[18]. Note that the move from the particle swarm at a period $t - 1$, $\left\{w_1^{t-1,i}, s_0^{t-1,i}\right\}_{i=1}^{N}$, to the conditional probabilities describing the likelihood function, $p\left(y_t | W_1^t, S_0, y^{t-1}; \gamma\right)$, requires the propagation of the states and the matching to observable variables according to equations 1 and 2. The final value of the conditional probability is obtained from the ability to evaluate the probability density function associated to measurement errors, $p\left(V_t; \gamma\right)$. Figure 2 shows the evolution of density for one state variable in the theoretical model – capital – and a red line with the time series of the observable over the sample simulation[5].

Another important issue regarding the algorithm is the initialization step. Despite the relevance of the topic, associated with many difficulties to obtain a good guess for the initial values of the swarm, particle filters here were initialized using the true set of states in every simulation – thus, S_0 is assumed to be known to the researcher during simulations. This assumption greatly simplifies the procedure and highlights the behavior of the filter excluding additional noise associated to the choice of an initial value.

The basic idea of the Bootstrap Filter is that particles that do not contribute to characterize the state vector in each point in time must be eliminated in favor of those with large weight in the distribution. As a consequence, the filter concentrates particles around the relevant regions of the state space of the model. From this perspective, the Sampling step in the algorithm, based in a small modification from the baseline Importance Sampling method (or, more specifically, the Sequential Importance Sampling method), is the core of the filter: with the basic Importance Sampling, each particle generated in the prediction step would be sampled with equal probabilities. However, it is well known that, for $t \rightarrow \infty$, there is a degeneracy problem such that, except for one specific particle, all the importance weights converge to zero. Furthermore, even the one particle with mass equal to one does not, necessarily, provide a good characterization of the state vector[6].

[5] Density is more concentrated in the initial simulations because the Bootstrap filter was initialized, on purpose, at the true state S_0.

[6] See Robert and Casella (2002)[26], chapter 14.

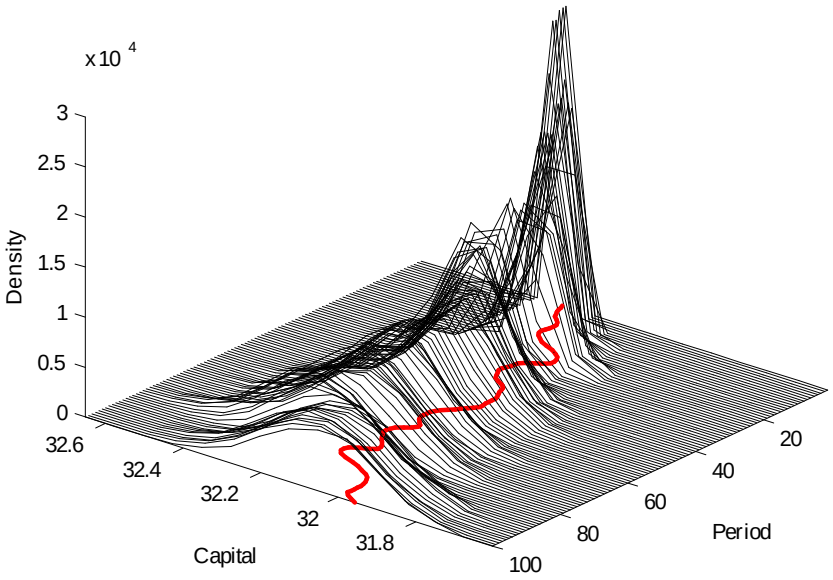

Figure 2: Probability Density Function of Capital – 100 Periods, $N = 30,000$

The degeneracy problem described above imposes, implicitly, a lower bound for the number of particles N describing the system. It is also worth noting that the main theorems regarding the convergence of the empirical distributions generated by Bootstrap Filter to their true values are only valid as $N \to \infty$. The main reason for that is the nature of the Importance Sampling mechanism, as the estimate of the objective function $p\left(W_1^t, S_0 | y^{t-1}; \gamma\right)$ is based on the ratio of two estimates – thus, a biased approximation of the true values for a fixed N. An almost-sure convergence based in the Strong Law of Large Numbers holds as N increases. To be more precise, Künsch (2005)[22] shows that the number of particles must increase at an exponential rate to ensure convergence in total variation of the filter. Interestingly, the result holds irrespective of the sample size[7]. The result in Künsch (2005)[22] is particularly relevant, as it stresses the trade-off between the accuracy (or convergence) of the Bootstrap Filter and the required time to compute a large number of particles – it is the analytical proof of the argument made at the introduction with figure 1.

One method to evaluate the quality of the approximation of the density of states in particle filters is the computation of the so-called effective number of particles of the swarm[8]. The main idea is to use the variance of the importance weights as a tool to evaluate the quality of the distribution. The number of effective particles characterizing the observations in period t in a

[7] Other results of convergence of particle filters are in Crisan and Doucet (2002)[10].
[8] See Liu and Chen (1995)[23].

particle filter with swarm of size N is estimated as:

$$N^{eff} = \frac{1}{\sum_{i=1}^{N} \left(q_t^i\right)^2}$$

There are other methods to generate proposals of the state vector to the next period. Pitt and Shephard (1999)[25], in particular, note that the presence of outliers in the data usually requires larger values of N in the Bootstrap Filter in order to obtain a good approximation of the density in that step, as the importance weights become heavily skewed towards a small number of sample draws. They propose a filter where the importance weights are computed using an auxiliary function including not only past, but also current values of the observable variables, Y^T, minimizing, thus, the effect of outliers in the distribution of weights.

The algorithm of the Auxiliary Particle Filter closely follows the last steps of the Bootstrap Filter: it adds a preliminary approximation of the state variables before starting the prediction step described in the Bootstrap Filter. The first step consists of a propagation of states and evaluation of densities based on the functional form of the state space representation, except for the distribution of shocks. With the initial step, information about the distribution of the densities at time t is included in the prediction and filtering steps of the Bootstrap Filter, at the cost of an extra computational time and some Monte Carlo noise from the additional resampling step. The algorithm below summarizes the Auxiliary Particle Filter:

Algorithm 2: Auxiliary Particle Filter

1) Initialization: $t = 0$. For $i = 1, ..., N$, initialize $p\left(W_1^{t-1}, S_0 | y^{t-1}; \gamma\right)$ and set $t = 1$.

2) Auxiliary variable: Generate N proposals of $\left\{\widetilde{s}_0^{t-1,i}\right\}_{i=1}^{N}$ from the distribution $p\left(W_1^{t-1}, S_0 | y^{t-1}; \gamma\right)$.

3) Auxiliary filtering/sampling: Evaluate importance weights $\widetilde{q}_t^i$ and resample with replacement N draws from $\left\{\widetilde{s}_0^{t|t-1,i}\right\}_{i=1}^{N}$ based on probabilities $\widetilde{q}_t^i$.

4) Prediction: Generate N proposals of $\left\{w_1^{t-1,i}, \widetilde{s}_0^{t|t-1,i}\right\}_{i=1}^{N}$ from the distribution $p\left(W_1^{t|t-1}, \widetilde{S}_0 | y^{t-1}; \gamma\right) = p\left(W_{1,t}; \gamma\right) p\left(W_1^{t-1}, \widetilde{S}_0 | y^{t-1}; \gamma\right)$.

5) Filtering: Evaluate importance weights q_t^i.

6) Sampling: Resample with replacement N draws from $\left\{w_1^{t|t-1,i}, \widetilde{s}_0^{t|t-1,i}\right\}_{i=1}^{N}$ based on probabilities q_t^i. Call the draws $\left\{w_1^{t,i}, s_0^{t,i}\right\}_{i=1}^{N}$. If $t < T$, set $t \rightsquigarrow t + 1$ and go to step 2. Stop otherwise.

Note from the description of algorithm 2 that steps 4, 5 and 6 are exactly the same as steps 2, 3 and 4 of the Bootstrap Filter, except for the use of the distribution $\widetilde{S}_0$, built based on the auxiliary steps, and the definition of q_t^i, now incorporating information about states in t on prediction. The use of the state equation as the auxiliary variable before the filtering step has one major problem, associated to the presence of outliers in the state variables: if the volatility

of structural shocks – described in matrix W_t – is too high, the value of $E\left[S_t|S_{t-1}\right]$ is not a good approximation to $p\left(W_1^{t|t-1}, S_0|y^{t-1};\gamma\right)$. As Arulampalam et al. (2002)[5] point out, a large variance of structural shocks might even deteriorate the performance of the Auxiliary Particle Filter.

In a note on the computational requirements associated with particle filters, all computations were made in a personal computer with Intel Core i7 processor and 8Gb of RAM, with all programs compiled in C++ language. Significant reductions of the computational burden were obtained through the use of a parallel distribution of tasks associated to the prediction step in the algorithm. OpenCL code was written to distribute the tasks to a NVIDIA GeForce GTX 570 graphics processing unit (GPU), easily integrating the parallel tasks to the main C++ code. Even with the use of OpenCL, the prediction step, including the propagation of states and computation of $p\left(y_t|w_1^{t|t-1,i}, s_0^{t|t-1,i}, y^{t-1};\gamma\right)$, represented between 30 and 72% of the time spent in the evaluation of the likelihood, depending on the number of particles used. The time presented in figure 1 and in the next sections is obviously a simple benchmark for particle filters accelerated by GPU computation, as only trivial tasks were moved to parallel computation in GPU and a preprocessed package – Dynare++ – was used to compute the solution of the model. Other steps of the filter, mainly the sampling step, could also be run in parallel, as proposed in Bolic, Djuric and Hong (2005)[8]. However, the sampling step, even applied at every step of the algorithm with serial processing, represented only 2.2% of the time used to compute the likelihood with 50,000 particles in the Bootstrap Filter

The next section describes the theoretical model used for simulations, its calibration, solution method and a detailed description of the observation equations.

3 Theoretical Model

The model used for analysis is based on a version of Christiano, Eichenbaum and Evans' (2005)[9] model presented in Schmitt-Grohé and Uribe (2004)[27]. The model contains the most common features of a DSGE model in terms of real and nominal rigidities: sticky prices and wages, demand for money from households justified by its presence in the utility function, firms' demand for money in the form of a cash-in-advance constraint, investment adjustment costs, variable capacity utilization of capital, internal habit formation in consumption and monopolistic competition in the market of goods and labor. In terms of driving forces, the model is characterized by three shocks: a stationary productivity shock affecting the production function, government spending shocks characterizing exogenous shifts in demand and a monetary policy shock. Relative to the original Christiano, Eichenbaum and Evans (2005)[9] model, Schmitt-Grohé and Uribe (2004)[27] transformed the monetary policy function to have the monetary authority following a Taylor-type policy rule to set interest rates, instead of a money supply function[9]. The authors also

[9] The Taylor-type policy rule was used in Christiano, Eichenbaum and Evans (2005)[9] as a robustness test. The estimation procedure, however, was conducted on the model with the money supply function.

included the approach of Yun (1996)[30] to characterize the distortions in allocations generated from sticky wages and prices in the economy, well suited for the use of non-linear methods to solve the model in a steady state where inflation is not zero.

From a computational perspective, the model not only reflects the current size of most of models used in the literature for empirical work, with a total of 13 state variables including three structural shocks, but it also provides the necessary flexibility to reduce the number of such state variables with small modifications in parameters. A modified version of the model, with different number of state variables and shocks, will be used to evaluate the sensitivity of the number of particles in the filter with respect to the size of the model. In this section, a brief presentation highlights the main equilibrium conditions of the economy and emphasizes the main parameters used to reset the number of state variables of the model. Original sources are highly recommended for further details of the model. Notation is kept as close as possible to Schmitt-Grohé and Uribe (2004)[27].

There is a large population of identical agents with utility function $U\left(c_t, h_t, m_t^h\right)$ characterized by the presence of internal habit persistence in consumption and a demand for real money balances. Agents choose in each period the optimal bundle of consumption from a continuum of differentiated goods indexed by i and aggregated using a Dixit-Stiglitz function. In each period, agents minimize the expenditure in consumption for each of the i goods, given its prices. Agents also supply labor to a central union able to negotiate wages in each of the j segments of the labor market in the economy. Given aggregate wages and the overall demand of labor of the economy, the union supplies all labor necessary to satisfy the demand in each market by adjusting nominal wages W_t^j. In terms of intertemporal decisions, agents choose the expenditure in investment for next period's capital, i_t, and the utilization of capital in the current period, u_t. Investment is subject to adjustment costs and agents solve a cost-minimization problem in each period to determine the amount of good i that is used in investment. Capital is rented to firms at the rate r_t^k. Still related to the intertemporal decisions of agents, there is a complete set of nominal state-contingent assets in the economy characterized by the stochastic nominal discount factor between period t and period $t+1$, $r_{t,t+1}$.The presence of complete markets imply that the nominal interest rate on a one-period bond, R_t, must be equal to the reciprocal of $r_{t,t+1}$. Define λ_t the Lagrange multiplier for the agents' budget constraint and $\lambda_t q_t$ the Lagrange multiplier for the equation describing capital accumulation. The main equilibrium conditions for households are given by:

$$U_c\left(c_t, h_t, m_t^h\right) - b\beta U_c\left(c_{t+1}, h_{t+1}, m_{t+1}^h\right) = \lambda_t \tag{3}$$

$$\lambda_t r_{t,t+1} = \beta E_t \frac{\lambda_{t+1}}{\pi_{t+1}} \tag{4}$$

$$-U_h\left(c_t, h_t, m_t^h\right) = \frac{\lambda_t w_t}{\widetilde{\mu}_t} \tag{5}$$

$$\lambda_t q_t = \beta E_t \lambda_{t+1}\left[r_{t+1}^k u_{t+1} - a\left(u_{t+1}\right) + q_{t+1}\left(1-\delta\right)\right] \tag{6}$$

$$\lambda_t = \lambda_t q_t \left[1 - S\left(\frac{i_t}{i_{t-1}}\right) - \left(\frac{i_t}{i_{t-1}}\right) S'\left(\frac{i_t}{i_{t-1}}\right)\right] + \beta E_t \lambda_{t+1} q_{t+1} \left(\frac{i_{t+1}}{i_t}\right)^2 S'\left(\frac{i_{t+1}}{i_t}\right) \quad (7)$$

$$r_t^k = a'(u_t) \quad (8)$$

$$\lambda_t = U_m\left(c_t, h_t, m_t^h\right) + \beta E_t \frac{\lambda_{t+1}}{\pi_{t+1}} \quad (9)$$

$$k_{t+1} = (1-\delta) k_t + i_t \left(1 - S\left(\frac{i_t}{i_{t-1}}\right)\right) \quad (10)$$

$$r_{t,t+1} R_t = 1 \quad (11)$$

In these equations, $a(u_t)$ is the cost for changing the utilization of capital in each period, $S\left(\frac{i_t}{i_{t-1}}\right)$ is the adjustment cost of investment, $\pi_t = \frac{P_t}{P_{t-1}}$ is inflation, $w_t = \frac{W_t}{P_t}$ is the average real wage across the j sectors of the economy and U_c, U_h and U_m are the first order derivatives of the utility function in terms of consumption, labor supply and money, respectively. The parameter characterizing the degree of habit persistence in consumption is given by b, while δ is the depreciation rate.

The intratemporal condition relating the disutility of labor and real wages shows an extra term given by $\widetilde{\mu}_t$ originated at the Lagrange multiplier of the constraint relating the demand for labor in sector j and the aggregate demand for labor: $h_t = h_t^d \int_0^1 \left(\frac{w_t^j}{w_t}\right)^{-\widetilde{\eta}} dj$. This wedge in the equilibrium condition is a result of sticky wages, as the marginal rate of substitution between consumption and labor is always equal to the real wage only in a world with flexible wages. The sticky wage mechanism takes the form of a Calvo-style function, with a share of $1 - \widetilde{\alpha}$ of the labor market where the union can set wages optimally in period t. The share of labor market not allowed to optimally adjust nominal wages in the current period follows the indexation rule given in market j by $W_t^j = W_{t-1}^j \pi_{t-1}^{\widetilde{\chi}}$, for $\widetilde{\chi} \in [0\ 1]$. The presence of Calvo stickiness in the model generates the usual equilibrium condition to determine real wages in each period:

$$E_t \sum_{s=0}^{\infty} \widetilde{\alpha}^s \lambda_t \beta^s \left(\frac{\widetilde{w}_t}{w_{t+s}}\right)^{-\widetilde{\eta}} h_{t+s}^d \prod_{k=1}^{s} \left(\frac{\pi_{t+k}}{\pi_{t+k-1}^{\widetilde{\chi}}}\right)^{\widetilde{\eta}} \left[\frac{\widetilde{\eta}-1}{\widetilde{\eta}} \frac{\widetilde{w}_t}{\prod_{k=1}^{s} \left(\frac{\pi_{t+k}}{\pi_{t+k-1}^{\widetilde{\chi}}}\right)} - \frac{w_{t+s}}{\widetilde{\mu}_{t+s}}\right] = 0$$

where $\widetilde{w}_t$ is the optimal wage defined in markets that were allowed to do so. At this point, Schmitt-Grohé and Uribe (2004)[27] define two auxiliary variables, f_t^1 and f_t^2, to split the equilibrium condition and rewrite the equation in a recursive form given by:

$$f_t^1 = \lambda_t \left(\frac{\widetilde{w}_t}{w_t}\right)^{-\widetilde{\eta}} h_t^d + \widetilde{\alpha}\beta E_t \left(\frac{\pi_{t+1}}{\pi_t^{\widetilde{\chi}}}\right)^{\widetilde{\eta}-1} \left(\frac{\widetilde{w}_{t+1}}{\widetilde{w}_t}\right)^{\widetilde{\eta}} f_{t+1}^1 \quad (12)$$

$$f_t^2 = \frac{\lambda_t}{\widetilde{\mu}_t} w_t \left(\frac{\widetilde{w}_t}{w_t}\right)^{-\widetilde{\eta}} h_t^d + \widetilde{\alpha}\beta E_t \left(\frac{\pi_{t+1}}{\pi_t^{\widetilde{\chi}}}\right)^{\widetilde{\eta}} \left(\frac{\widetilde{w}_{t+1}}{\widetilde{w}_t}\right)^{\widetilde{\eta}} f_{t+1}^2 \quad (13)$$

$$\frac{(\widetilde{\eta}-1)}{\widetilde{\eta}}\widetilde{w}_t f_t^1 \ = f_t^2 \tag{14}$$

Firms use capital and labor to produce a differentiated good i in order to satisfy the demand of consumption and investment from households and the government spending. Firms must finance a share ν of the labor costs in order to produce, characterizing the cash-in-advance constraint of the firm. The first order conditions of the cost minimization problem in each period of the firm, the firms's demand and the cash-in-advance constraint of the firm are given by:

$$w_t\left[1+\nu\frac{R_t-1}{R_t}\right] = mc_{it}z_t F_h\left(k_{it},h_{it}\right) \tag{15}$$

$$r_t^k = mc_{it}z_t F_k\left(k_{it},h_{it}\right) \tag{16}$$

$$z_t F\left(k_{it},h_{it}\right)-\psi = \left(\frac{P_{it}}{P_t}\right)^{-\eta}\left(c_t+i_t+g_t+a\left(u_t\right)k_t\right) \tag{17}$$

$$m_{it}^f \geq \nu w_t h_{it} \tag{18}$$

$$\log\left(z_t\right) = \alpha_z \log\left(z_{t-1}\right)+\varepsilon_t^z \qquad \varepsilon_t^z \sim N\left(0,\sigma_z\right) \tag{19}$$

In the equations above, $F\left(k_{it},h_{it}\right)$ is the production function, z_t is a technological shock, mc_{it} is the inverse of the markup over prices, derived from the Lagrange multiplier of the firms' problem with respect to the equation describing the demand for goods and m_{it}^f is the money demand of firm i. The technology for production is non-convex because of the presence of a fixed cost measured by ψ. The first term next to real wages in equation 15 characterizes the effect of the cash-in-advance constraint for firms and the extra-cost to obtain money to finance the wage bill.

Prices are set by firms following the basic Calvo model with indexation: a share $1-\alpha$ of firms is allowed in each period to optimally set their prices. Firms not allowed to optimally choose prices follow the indexation rule $P_{it}=P_{it-1}\pi_{t-1}^{\chi}$, as in the rule applied to the sticky wage process. If a firm is allowed to reoptimize prices, the new price $\widetilde{P}_t$, which is the same for all firms in the same situation in that period, characterizes the first-order condition of the maximization problem:

$$E_t\sum_{s=0}^{\infty} r_{t,t+s}P_{t+s}\alpha^s\left(\frac{\widetilde{P}_t}{P_t}\right)^{-\eta}\prod_{k=1}^{s}\left(\frac{\pi_{t+k-1}^{\chi}}{\pi_{t+k}}\right)^{-\eta} y_{t+s}\left[\frac{\eta-1}{\eta}\left(\frac{\widetilde{P}_t}{P_t}\right)\prod_{k=1}^{s}\left(\frac{\pi_{t+k-1}^{\chi}}{\pi_{t+k}}\right)-mc_{it+s}\right]=0$$

Using the same approach described for wage stickiness, define $\widetilde{p}_t=\widetilde{P}_t/P_t$ and rewrite the first-order condition in recursive terms, after defining two auxiliary variables x_t^1 and x_t^2:

$$x_t^1 = y_t mc_t\widetilde{p}_t^{-\eta-1}+\alpha\beta E_t\left(\frac{\lambda_{t+1}}{\lambda_t}\right)\left(\frac{\widetilde{p}_t}{\widetilde{p}_{t+1}}\right)^{-\eta-1}\left(\frac{\pi_t^{\chi}}{\pi_{t+1}}\right)^{-\eta}x_{t+1}^1 \tag{20}$$

$$x_t^2 = y_t\widetilde{p}_t^{-\eta}+\alpha\beta E_t\left(\frac{\lambda_{t+1}}{\lambda_t}\right)\left(\frac{\widetilde{p}_t}{\widetilde{p}_{t+1}}\right)^{-\eta}\left(\frac{\pi_t^{\chi}}{\pi_{t+1}}\right)^{1-\eta}x_{t+1}^2 \tag{21}$$

$$x_t^1 = \frac{\eta - 1}{\eta} x_t^2 \tag{22}$$

In the model, the government collects lump sum taxes and print money in order to finance an exogenous stream of expenditure. Money supply satisfies the demand from firms and households. The budget constraint of the government is given by:

$$g_t = \tau_t + m_t - \frac{m_{t-1}}{\pi_t} \tag{23}$$

$$m_t = m_t^f + m_t^h \qquad m_t^f = \int_0^1 m_{it}^f di \tag{24}$$

$$\log\left(\frac{g_t}{g}\right) = \alpha_g \log\left(\frac{g_{t-1}}{g}\right) + \varepsilon_t^g \qquad \varepsilon_t^g \sim N\left(0, \sigma_g\right) \tag{25}$$

Monetary policy is described by the Taylor rule described in Christiano, Eichenbaum and Evans (2005)[9] as an alternative formulation for the process characterizing interest rates:

$$\log\left(\frac{R_t}{R}\right) = \rho \log\left(\frac{R_{t-1}}{R}\right) + (1-\rho)\left(\alpha_\pi \log\left(\frac{\pi_t}{\pi}\right) + \alpha_y \log\left(\frac{y_t}{y}\right)\right) + \varepsilon_t^r \tag{26}$$

$$\varepsilon_t^r \sim N\left(0, \sigma_r\right)$$

In the equation, nominal interest rates R_t are determined based on deviations of inflation and output from their steady state values, an autoregressive term and an exogenous shock ε_t^r.

Finally, in terms of aggregation, it is important to characterize the evolution of aggregate prices and wages and the effect of Calvo stickiness in terms of price and wage dispersion in the cross-sectional dimension. First, note that the price and wage indexes of the economy change as a result of the share of firms and workers allowed to optimize prices. With little algebra, prices and wages are expressed by:

$$1 = \alpha\left(\frac{\pi_{t-1}^\chi}{\pi_t}\right)^{1-\eta} + (1-\alpha)\,\widetilde{p}_t^{1-\eta} \tag{27}$$

$$1 = \widetilde{\alpha}\left(\frac{w_{t-1}}{w_t}\right)^{1-\widetilde{\eta}}\left(\frac{\pi_{t-1}^{\widetilde{\chi}}}{\pi_t}\right)^{1-\widetilde{\eta}} + (1-\widetilde{\alpha})\left(\frac{\widetilde{w}_t}{w_t}\right)^{1-\widetilde{\eta}} \tag{28}$$

In order to characterize the effects of price dispersion, consider the relation between supply and demand for good i in equation 17. It is possible to prove that the capital-labor ratio is the same across firms. Thus, aggregating both sides of equation 17 and noting that $u_t k_t = \int_0^1 k_{it} di$ and $h_t^d = \int_0^1 h_{it} di$, the only term containing variables at the firm level is $s_t = \int_0^1 \left(\frac{P_{it}}{P_t}\right)^{-\eta} di$. It is possible to express s_t in recursive form as:

$$s_t = (1-\alpha)\,\widetilde{p}_t^{-\eta} + \alpha\left(\frac{\pi_t}{\pi_{t-1}^\chi}\right)^\eta s_{t-1} \tag{29}$$

Equilibrium in the labor market follows a similar pattern in terms of supply and demand:

$$h_t = \widetilde{s}_t h_t^d \tag{30}$$

$$\widetilde{s}_t = (1-\widetilde{\alpha}) \sum_{s=0}^{\infty} \widetilde{\alpha}^s \left(\frac{\widetilde{W}_{t-s}}{W_t} \prod_{k=1}^{s} \pi_{t+k-s-1}^{\widetilde{\chi}} \right)^{-\widetilde{\eta}}$$

The new variable $\widetilde{s}_t$, measuring the distortion generated in labor markets due to wage dispersion, can be expressed recursively as:

$$\widetilde{s}_t = (1-\widetilde{\alpha}) \left(\frac{\widetilde{w}_t}{w_t} \right)^{-\widetilde{\eta}} + \widetilde{\alpha} \left(\frac{w_{t-1}}{w_t} \right)^{-\widetilde{\eta}} \left(\frac{\pi_{t-1}^{\widetilde{\chi}}}{\pi_t} \right)^{\widetilde{\eta}} \widetilde{s}_{t-1} \tag{31}$$

The model contains three shocks, $\varepsilon_t^g, \varepsilon_t^{\widetilde{z}}$ and ε_t^r, two autoregressive processes, g_t and z_t and a total of 29 variables: $y_t, c_t, i_t, g_t, k_t, h_t, h_t^d, mc_t, \pi_t, \widetilde{p}_t, w_t, \widetilde{w}_t, r_t^k, R_t, s_t, \widetilde{s}_t, \tau_t, m_t, m_t^f, m_t^h, \lambda_t, x_t^1, x_t^2$, $z_t, f_t^1, f_t^2, \widetilde{\mu}_t, r_{t,t+1}, q_t$ and u_t. The system characterizing the economy is given by the set of equations 3 to 31. The characterization of a stationary equilibrium of the economy requires initial conditions for $c_{-1}, i_{-1}, g_{-1}, k_0, \pi_{-1}, w_{-1}, R_{-1}, s_{-1}, \widetilde{s}_{-1}, z_{-1}$, and exogenous processes for $\left\{\varepsilon_t^g, \varepsilon_t^{\widetilde{z}}, \varepsilon_t^r\right\}_{t=0}^{\infty}$.

3.1 Functional forms and calibration

In order to close the model, it is necessary to specify functional forms for the utility function, the production function, the investment adjustment cost and the rate of capital utilization. Sticking to the literature, functional forms are the same as in Christiano, Eichenbaum and Evans (2005)[9]:

$$U = \log\left(c_t - bc_{t-1}\right) - \frac{\phi_0}{2} h_t^2 + \phi_1 \frac{\left(m_t^h\right)^{1-\sigma_m}}{1-\sigma_m}$$

$$F\left(k,h\right) = k^{\theta} h^{1-\theta}$$

$$S\left(\frac{i_t}{i_{t-1}}\right) = \frac{\kappa}{2}\left(\frac{i_t}{i_{t-1}} - 1\right)^2$$

$$a\left(u\right) = \gamma_1\left(u-1\right) + \frac{\gamma_2}{2}\left(u-1\right)^2$$

The baseline calibration is composed by the same parameters describing the structural model in Christiano, Eichenbaum and Evans (2005)[9]. It assumes quarterly frequency of data, an annual real interest rate of 3% per year, capital share of 36% of value added and an annual depreciation rate of 10%. There is full indexation of prices and wages to the past inflation and long run inflation is set at 4.2% per year. It also assumes that households hold 44% of real money holdings and firms must finance 100% of the wage bill. The Taylor rule is characterized by values given in Christiano, Eichenbaum and Evans (2005)[9] for the US. The calibration of shocks assumes the same first-order serial correlation for government spending and productivity used

in Schmitt-Grohé and Uribe (2004)[27]. Parameters describing the volatility of shocks, σ_z and σ_g, also follows Schmitt-Grohé and Uribe (2004)[27], matching moments estimated for the Solow residual and the government purchases estimated in the literature. In order to match output volatility in the model, a scale parameter is used to adjust these two values, thus preserving the ratio of the volatilities. The volatility of the monetary policy shock, σ_r, is calibrated to match the share of output's variance explained by the shock 20 quarters ahead, as estimated in Christiano, Eichenbaum and Evans (2005)[9] for the US. Table 1 shows the calibrated values.

Table 1: Calibration — CEE (2005)

Parameter		Parameter		Parameter	
β	$1.03^{-0.25}$	κ	2.48	ρ	0.8
θ	0.36	b	0.65	α_π	1.5
δ	0.025	ϕ_1	0.539	α_y	0.1
α	0.60	σ_m	10.62	α_g	0.96
$\widetilde{\alpha}$	0.64	γ_2/γ_1	0.01	σ_g	0.02
η	6	ν	1	α_z	0.979
$\widetilde{\eta}$	21	χ	1	σ_z	0.0072
π	$1.042^{0.25}$	$\widetilde{\chi}$	1	σ_r	0.0024
		"Big Ratios"			
u	1	Profits share	0	g/y	0.18

As discussed in the introduction, the model is flexible enough to accommodate simpler formulations using small changes in the baseline calibration. The alternative model used for testing robustness of results has flexible wages ($\widetilde{\alpha} = 0$), no habit persistence in consumption ($b = 0$), no adjustment cost in investment ($\kappa = 0$), no inflation indexation ($\chi = 0$) and the Taylor rule does not have an autoregressive component ($\rho = 0$). In terms of the definition of a competitive equilibrium, these modifications result in a requirement of initial conditions for only four state variables: $g_{-1}, k_0, s_{-1}, z_{-1}$.

3.2 Solution method, state vector and observable variables

The non-linear nature of the model imposes the use of a solution method capable of numerically approximating the policy functions. Here, the model is solved using a second-order perturbation[10] due to the good properties in terms of accuracy reported in Aruoba, Fernández-Villaverde and Rubio-Ramírez (2006)[6]. Perturbation approximates the policy functions of the model in terms of Taylor expansions around the deterministic steady state of the state variables, S_t and the vector of shocks W_t defined in equations 1 and 2. Set a vector stacking both vectors of states and structural shocks $Z_t = [(S_t - S); W_t]$. For a given variable Y_t, the second-order perturbation

[10] See Judd (1998)[20] for an introduction to perturbation methods. Solutions of the models were obtained using the package Dynare++, due to the use of symbolic computation to obtain the coefficients of the approximations. For details on Dynare++, please refer to Kamenik (2011)[21].

results in the following functional form for the policy function:

$$Y(Z_t) = Y + G_0 + G_1 Z_t + G_2(Z_t \otimes Z_t)$$

From the solution, Y is the deterministic steady state value of the variable Y_t, G_0 is a matrix adjusting for the steady state of the volatility of shocks and matrices G_1 and G_2 are functions of the derivatives of Y with respect to the vector Z. Note, particularly, that second-order perturbations highlight the role of uncertainty in the economy described above, as the solution is not certainty-equivalent as in the case of linear approximations of the policy functions.

In order to compute the likelihood of the model, simulated time series of output, consumption, investment, nominal interest rates and inflation were generated for the general model – with the complete set of states – and used throughout the exercises. Time series were simulated for 200 periods starting at the deterministic steady state, keeping only the last 100 observations. Additive independent Gaussian measurement errors were included to set the observable variables. The volatility of the measurement errors is arbitrarily set to the same value of the ergodic volatility of each variable[11]. Observation equations describing the model are linear with respect to the variables:

$$\begin{aligned}
y_t^{obs} &= y_t + \varepsilon_t^y, & \varepsilon_t^y &\sim N\left(0, \sigma_y^{obs}\right) \\
c_t^{obs} &= c_t + \varepsilon_t^c, & \varepsilon_t^c &\sim N\left(0, \sigma_c^{obs}\right) \\
i_t^{obs} &= i_t + \varepsilon_t^i, & \varepsilon_t^i &\sim N\left(0, \sigma_i^{obs}\right) \\
R_t^{obs} &= R_t + \varepsilon_t^R, & \varepsilon_t^R &\sim N\left(0, \sigma_R^{obs}\right) \\
\pi_t^{obs} &= \pi_t + \varepsilon_t^\pi, & \varepsilon_t^\pi &\sim N\left(0, \sigma_\pi^{obs}\right)
\end{aligned}$$

4 The Size of the Swarm and Filters' Performance

The main characteristics of the two particle filters are evaluated using a small Monte Carlo exercise, where the likelihood of the model is simulated 1000 times over 20 artificial samples of 100 periods. The same set of 20 samples is used in each exercise, changing only the number of particles characterizing the particle filter. Each time series was simulated for 200 periods, with the first half of the sample discarded in order to minimize the effects of initial values, reaching, thus, the target of 100 observations.

The first exercise to evaluate the performance of the Bootstrap and the Auxiliary Particle Filter computes an estimate of the number of effective particles and the ratio between the standard-deviation and the mean of the simulated log-likelihood. The last statistic was the base

[11] The use of noisy observations helps simulations in two dimensions: first, it provides a more realistic framework compared to empirical work using real data; second, the high volatility of measurement errors stabilizes the evaluation of densities in the particle filter, as the covariance matrix of the measurement errors does not get too close to singularity.

for Fernández-Villaverde and Rubio-Ramírez (2004)[12] exercise on setting the appropriate number of particles in the Bootstrap Filter. Results comparing both filters are presented in table 2. Under the baseline framework, the Bootstrap Filter presents a gap between the actual and the effective number of particles estimated around 42%, irrespective of the size of the particle swarm. The two main thresholds used in the literature to evaluate the filter's quality, based on the ratio between effective and actual number of samples, are violated in a significant number of cases: 60% of the total number of simulations moves below the first threshold ($N^{eff} = 2N/3$) and almost one third of simulations have an effective number of particles below $N/2$. It is worth noting that the increase in the size of the swarm clearly affects the tails of the probability density function of states, as the smallest number of effective particles found in simulations grows at a higher rate than the average number of effective particles.

Table 2: Effective Particles – Bootstrap Filter and Auxiliary Particle Filter

Samples		$\widehat{N^{eff}}$	$\min\left(\widehat{N^{eff}}\right)$	$\widehat{N^{eff}} \leq \frac{2N}{3}(\%)$	$\widehat{N^{eff}} \leq \frac{N}{2}(\%)$	$\frac{Std(L)}{\mu(L)}(\%)$
1000	B.Filter	579.52	3.13	59.94	32.87	0.047
	Aux.Filter	690.97	37.99	36.52	13.60	0.056
5000	B.Filter	2892.8	19.55	60.29	32.90	0.021
	Aux.Filter	3448.2	333.83	36.63	13.83	0.026
10000	B.Filter	5784.0	47.28	60.36	32.91	0.015
	Aux.Filter	6893.8	562.35	36.67	13.88	0.018
20000	B.Filter	11566	113.48	60.39	32.92	0.011
	Aux.Filter	13784	1203.08	36.71	13.91	0.013
30000	B.Filter	17349	257.08	60.41	32.94	0.009
	Aux.Filter	20674	1650.5	36.72	13.92	0.011
40000	B.Filter	23131	333.93	60.41	32.93	0.007
	Aux.Filter	27563	1976.1	36.73	13.92	0.010
50000	B.Filter	28913	714.03	60.42	32.94	0.007
	Aux.Filter	34453	1964.8	36.71	13.92	0.009

Table 2 also shows the improvement in terms of effective particles of using the Auxiliary Particle Filter. There is a significant gain on the gap between the actual and the effective size of the particle swarm, reducing the number of lost particles from 42% in the Bootstrap Filter to 31% in the Auxiliary Particle Filter. Improvements are also verified at the tails of the probability density function, specially when the size of the swarm is small. The minimum value of the number of effective particles is more than ten times larger in the Auxiliary Particle Filter when compared to the Bootstrap filter when N is set between 1,000 and 20,000 particles. The improvement in the characterization of the tails of the probability density function also reflects in the number of times the estimated size of the effective swarm violates the thresholds of two-thirds (36.7%) and half the size of the actual swarm (only 13.9%).

In terms of accuracy, the ratio between the standard-deviation and the mean of the log-likelihood, measuring the Monte Carlo variation of the filter, shows that the Bootstrap Filter is always more precise than the Auxiliary Particle Filter. However, the largest difference across

filters is smaller than 0.01%. One reason for the result is related to the resampling procedure applied twice for every evaluation of the density in the Auxiliary Particle Filter, which is known to increase the variation of Monte Carlo simulations. Some authors propose using the resampling step only at the end of the algorithm, skipping step 3 in algorithm 2.

An alternative way to compare the two particle filters is proposed in Smith (2011)[29]. The main idea is to compare the time necessary to generate one effective particle during the evaluation of the likelihood. Table 3 shows the average time, measured in seconds, to compute the log-likelihood and the average number of effective samples for each particle filter. The last column is ratio of average seconds per effective sample measured in the Auxiliary Particle Filter over the measurement of the Bootstrap Filter. This measure considers the cost of using a more demanding filter, in terms of computational requirements, adjusted by the gains expressed as a better characterization of the probability density function of states.

Table 3: Average Seconds per Effective Particles

Samples	Bootstrap Filter		Auxiliary Particle Filter		
	Seconds per Sample (1)	$\widehat{N^{eff}}$ (2)	Seconds per Sample (3)	$\widehat{N^{eff}}$ (4)	Speedup [(3)(2)]/[(4)(1)]
1000	1.2211	579.52	1.7294	690.97	1.1878
5000	1.3524	2892.8	1.9670	3448.2	1.2202
10000	1.7493	5784.0	2.8558	6893.8	1.3697
20000	2.3210	11566	4.0242	13784	1.4549
30000	2.8764	17349	5.4048	20674	1.5768
40000	3.6333	23131	7.0359	27563	1.6251
50000	4.1767	28913	8.3208	34453	1.6719

Table 3 shows that the use of the Auxiliary Particle Filter results in significant gains relative to the Bootstrap Filter, despite the large computational requirements associated with the former. With 50,000 particles in the swarm, as an example, while the average time to compute the likelihood almost doubles, the increase in the number of effective particles results in a loss of only 67%. For the smallest size of the swarm tested, time to compute the log-likelihood increases by 41%, while the loss adjusted for effective particles is only 18%.

Finally, it is worth establishing a minimum size of the particle swarm. Boers (1999)[7] proposes tracking the errors in model's state variables. Despite the significant number of variables that are not observable to the researcher, tracking the ability of the particle filter to match state variables in a DSGE model using artificial data can provide a benchmark in terms of the smallest number of particles necessary to properly describe the density function. For dataset d, the mean absolute percentage error of state variable s from the vector of state variables Z_t simulated in N particles across all periods t is defined as:

$$MAE_{d,s} = \frac{1}{DT}\sum_{t=1}^{T}\sum_{d=1}^{D}\left|\frac{\widehat{Z_t^{s,d}} - Z_t^{s,d}}{Z_t^{s,d}}\right|$$

where $\widehat{Z_t^{s,d}}$ is the simulated mean state vector Z_t approximated by N particles.

Figure 3 shows the mean absolute percentage error for each state variable of the model derived from each particle filter. For the sake of clarity, MAE values from the Bootstrap Filter are presented in the left axis, while values from the Auxiliary Particle Filter are shown in the right axis. First, note that, as expected, when comparing the scale of the two axis that the MAE for the Auxiliary Particle Filter is, most of the time, smaller than the MAE of the Bootstrap Filter. More specifically, the Bootstrap Filter only presents a better performance for productivity, government spending and price dispersion, with similar MAE for capital when the size of the swarm is large. The use of information on the current period, added by the extra two steps on the Auxiliary Particle Filter, reduces the dispersion of samples around the true value of state variables.

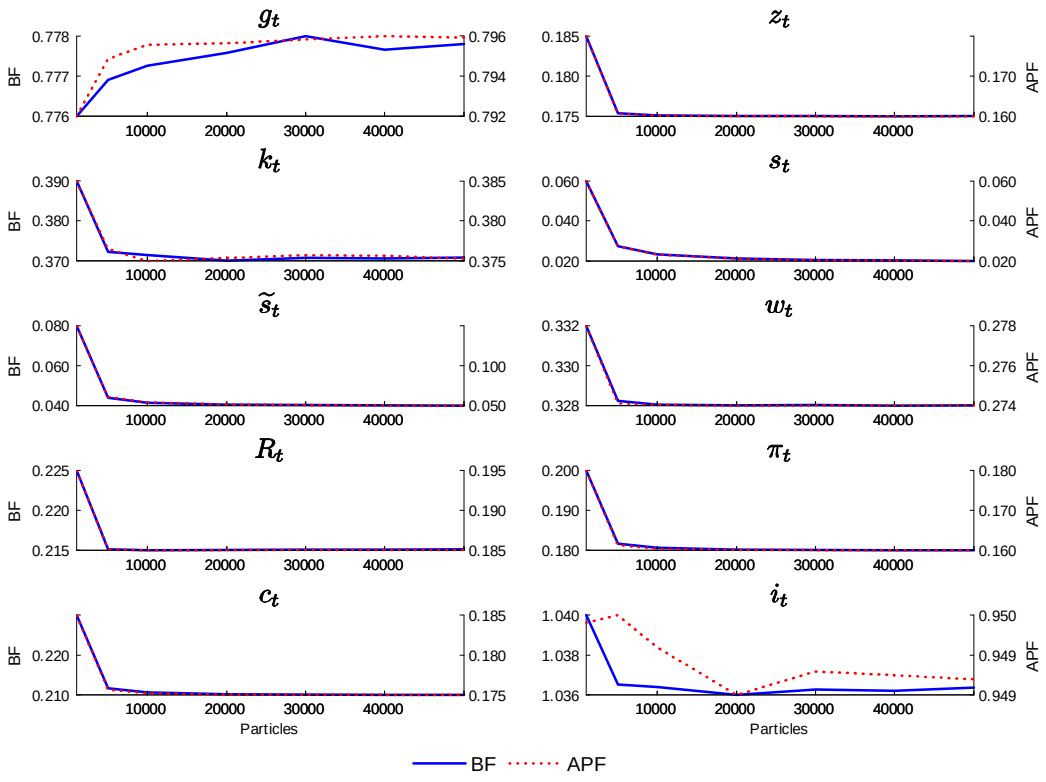

Figure 3: Mean Absolute Error (%) – State Variables

The second relevant observation is that MAE values seems to stabilize for values of N larger than 20,000. For all state variables, there is a large volatility in MAE values when the number of particles is small. However, for swarms with more than 20,000 particles, MAE fluctuations are very small, implying that moving to larger number of particles does not significantly changes

the quality of the approximation of density of states. Also, the MAE of some state variables – notably investment and capital – computed by the Auxiliary Particle Filter does not decrease when the size of the swarm is large – it only becomes less volatile. On the other hand, there is a clear downward path for MAE in the Bootstrap Filter as N increases.

4.1 Sensitivity to the Number of Time Series

The strategy here consists on testing both filters after removing one observable variable at a time, evaluating the number of effective particles, the speedup between the Bootstrap and the Auxiliary Particle Filter and the errors in tracking state variables. Observable variables are removed in the following sequence: inflation, interest rates, investment and consumption. Thus, the simulation using only one observable variable has output as the single time series in the observation equation. Table 4 shows the results of simulations in terms of mean of estimated effective particles for each filter, the speedup of using the Auxiliary Particle Filter over the Bootstrap Filter. The last column shows the standard deviation of the measurement error of the variable removed for the simulation. Thus, the value presented in the last column when the number of observations is equal to 4 is the standard deviation of the measurement error of inflation; the value when the number of observations is equal to 3 is the standard deviation of the measurement error of interest rates, and so on.

Table 4: Effective Particles and Number of Observables

Sample	Obs. Vars.	$\widehat{N^{eff}}(BF)$	$\widehat{N^{eff}}(APF)$	Speedup	σ_i^{obs} – Last Var.
1000	5	579.52	690.97	1.1878	—
	4	675.00	769.10	1.1842	0.0039
	3	790.19	867.48	1.2514	0.0040
	2	796.85	870.88	1.2501	0.0765
	1	828.03	889.84	1.2602	0.0159
20000	5	11566	13784	1.4549	—
	4	13493	15384	1.5237	0.0039
	3	15801	17347	1.5683	0.0040
	2	15933	17415	1.5903	0.0765
	1	16556	17791	1.6376	0.0159
50000	5	28913	34453	1.6719	—
	4	33731	38457	1.7141	0.0039
	3	39502	43364	1.8241	0.0040
	2	39833	43534	1.7658	0.0765
	1	41391	44474	1.8393	0.0159

There are two striking results on table 4. First, irrespective of the size of the particle swarm, removing observable variables of the system increases the cost of using the Auxiliary Particle

Filter, measured by the increase in the cost in seconds to generate one effective particle. Second, the improvements in the number of effective particles is directly correlated with the standard deviation of the measurement errors: the smallest is the volatility of the measurement error associated with the removed variable, the largest is the change in the estimated number of effective particles. As an example, notice that, when inflation and interest rates are removed from the estimation (model with 4 and 3 observable variables), the number of effective particles in the filter increases around 17% in the Bootstrap Filter and 12% in the Auxiliary Particle Filter, compared to the model with one more observable. When investment and consumption are removed (model with 2 and 1 observable variable), the increase in the number of effective particles is smaller than 4%.

The previous result is a consequence of the relation between the matrix of measurement errors and the estimated number of effective samples. The presence of Multivariate Gaussian and independent measurement errors implies that the inclusion of variables with a small volatility of measurement errors (i.e. standard error smaller than one) reduces the overall standard-error of the distribution. Thus, a larger number of samples drawn in the particle filter are lost in regions of negligible probability, reducing the estimated number of effective particles characterizing states. In order to avoid the problem of depletion from a tight distribution of measurement errors, An and Schorfheide (2007)[3] and An (2008)[2] fix the standard deviation of measurement errors to 20% of the standard error of observations.

Another consequence of changing the number of observable variables in the particle filter is the quality of tracking of state variables. Figure 4 shows the changes in MAE as a function of the number of observable variables for $N = 20,000$. For almost all states, there is a continuous reduction in the mean absolute error of state variables with the increase in the number of observable variables. This result is particularly impressive given the amount of noise added to observable variables as measurement error in the exercise.

Results in this section suggests that the number of time series used as observable variables can be one reason to choose the Auxiliary Particle Filter over the Bootstrap Filter in empirical applications. Reducing the number of observable variables favors the use of the Bootstrap Filter; however, it may also imply problems associated with the identification of structural parameters in the model in empirical applications.

4.2 Sensitivity to the Number of States

The objective of this section is to present results in a more realistic framework, as now the model used to compute the log-likelihood is not the same model generating artificial data. The model computing the likelihood is a special, simpler case of the baseline Christiano, Eichenbaum and Evans (2005)[9] model. It has flexible wages, no habit persistence in consumption, no adjustment cost in investment, no inflation indexation and no interest rate persistence in the Taylor rule, as discussed at the end of the calibration section. The simpler model has only four states (productivity, capital, price dispersion and government spending) and it is simulated with all 5

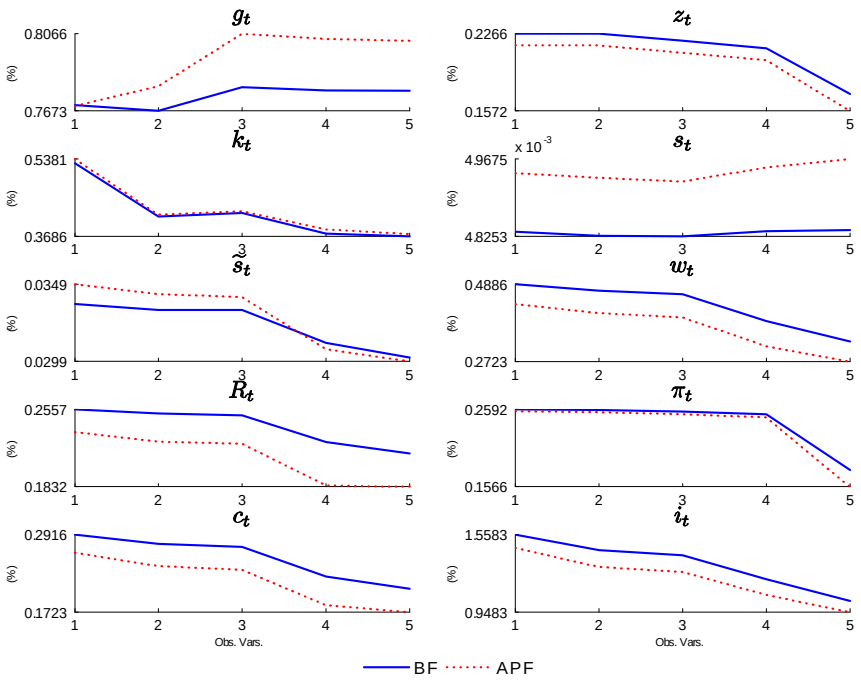

Figure 4: Mean Absolute Error and Number of Observables – $N = 20,000$

time series of the baseline exercise in the previous section. This exercise provides a more realistic framework since, for the econometrician, it is usually impossible to determine the exact data generating process, justifying the use of models to approximate relevant features of data.

Results on table 5 show that simpler models does not result in significant gains in terms of the number of effective particles, but increases the variability of the filter. The average number of effective particles is slightly higher than the estimated in simulations presented in table 2. The average loss of particles reduces in the Bootstrap Filter to 39%, which is a small gain when compared to the loss of 42% estimated in the baseline simulation. On the other hand, the Auxiliary Particle Filter shows a slightly larger loss of particles and larger variability of simulated likelihood. There is an average gap of 33% between the actual and the effective number of particles (31% in the complete model) and the Monte Carlo variation reaches 0.1% of the average value of the likelihood for a swarm of 1,000 particles. Based on this statistic, the precision of the Auxiliary Particle Filter is considerably smaller when compared to the Bootstrap Filter, as the Monte Carlo variation of the Auxiliary Particle Filter with a swarm of 50,000 particles is similar to the ratio of the Bootstrap Filter with only 10,000 particles.

In terms of time to obtain one effective particle, table 6 shows that the Auxiliary Particle Filter has a better performance when simulating the likelihood of a smaller model for a large number of particles. While the loss of using the Auxiliary Particle Filter reached 66% in the

Table 5: Effective Particles – Small Model

Samples		$\widehat{N^{eff}}$	$\min\left(\widehat{N^{eff}}\right)$	$\widehat{N^{eff}} \leq \frac{2N}{3}(\%)$	$\widehat{N^{eff}} \leq \frac{N}{2}(\%)$	$\frac{Std(L)}{\mu(L)}(\%)$
1000	B.Filter	609.97	2.50	52.96	28.55	0.037
	Aux.Filter	667.34	15.48	42.50	18.72	0.100
5000	B.Filter	3049.7	19.92	52.92	28.60	0.016
	Aux.Filter	3337.0	104.76	42.31	18.73	0.033
10000	B.Filter	6099.4	41.81	52.93	28.59	0.012
	Aux.Filter	6675.6	196.44	42.26	18.70	0.022
20000	B.Filter	12199	99.40	52.95	28.60	0.008
	Aux.Filter	13354	357.57	42.22	18.67	0.017
30000	B.Filter	18299	156.31	52.93	28.59	0.007
	Aux.Filter	20033	389.99	42.21	18.66	0.015
40000	B.Filter	24400	248.14	52.92	28.59	0.006
	Aux.Filter	26713	601.57	42.20	18.66	0.013
50000	B.Filter	30500	296.65	52.92	28.57	0.005
	Aux.Filter	33394	689.40	42.19	18.66	0.012

complete model, the smaller model resulted in a loss of almost 56% for a swarm of 50,000 particles. This result is important to highlight the gains of better computational resources when simulating the likelihood of these models: not only the average time to compute the likelihood is significantly reduced when the number of states is small, but also the effective loss of computing one effective particle is reduced, even for large values of N.

Table 6: Average Seconds per Effective Particles – Small Model

Samples	Bootstrap Filter		Auxiliary Particle Filter		
	Seconds per Sample (1)	$\widehat{N^{eff}}$ (2)	Seconds per Sample (3)	$\widehat{N^{eff}}$ (4)	Speedup [(3)(2)]/[(4)(1)]
1000	0.4693	609.97	0.5731	667.34	1.1162
5000	0.5066	3049.7	0.6472	3337.0	1.1676
10000	0.5742	6099.4	0.7736	6675.6	1.2310
20000	0.6811	12199	0.9896	13354	1.3272
30000	0.7683	18299	1.1929	20033	1.4183
40000	0.8844	24400	1.4542	26713	1.5018
50000	0.9955	30500	1.6986	33394	1.5583

Figure 5 shows the MAE for state variables of the small model. The simulation of the likelihood using a model that is different from the true data generating process results, as expected, in an increase of the error associated with tracking these variables. However, it is still surprising, when comparing with results in figure 4, that the increase in MAE is very small. The largest increase was on capital, where the MAE estimated changed from 0.37% with the complete model to 0.58% with the small model. There is also an important change with the MAE associated with price dispersion, but the average value in the small model is still very small – only 0.16%.

It is also worth noting that MAEs for capital and government spending are higher for the

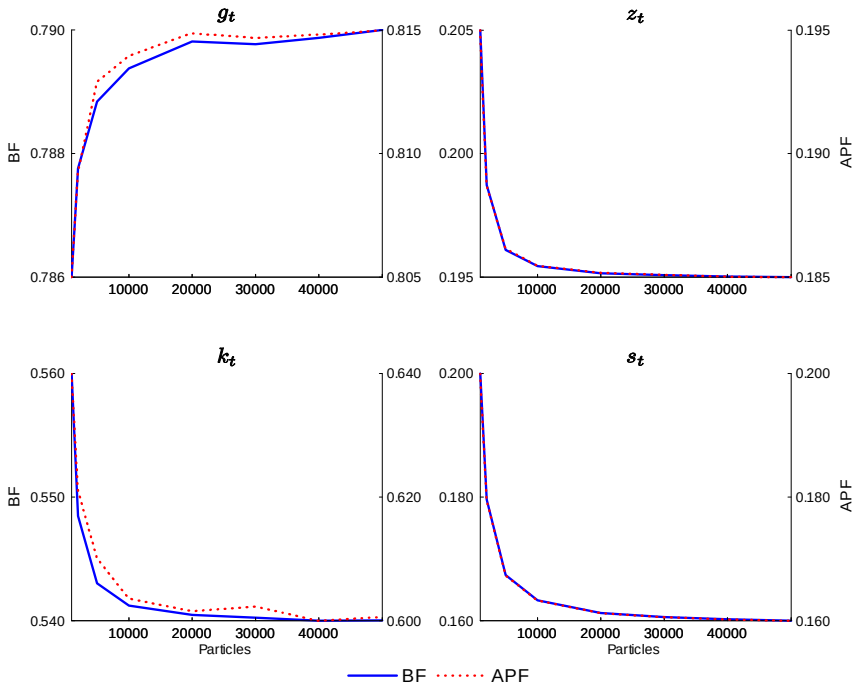

Figure 5: Mean Absolute Error (%) – State Variables – Small Model

Auxiliary Particle Filter, when compared to the outcome of the Bootstrap Filter. Despite appearing only for one state variable, it might be the case that the simple version of the model results in excessive volatility for capital, deteriorating the performance of the Auxiliary Particle Filter, as described in Arulampalam et al. (2002)[5].

Another feature presented on figure 5 is the relative stability of MAE for simulations with more than 20,000 particles in the swarm. Simulated MAE is very volatile when the particle swarm is small, with significant gains in MAE of capital and productivity for simulations with more than 10,000 particles. After 20,000 particles, the MAE for government spending, productivity and capital reach a stable level, especially for the Bootstrap Filter. Again, relating to results with the complete model, simulations with a swarm between 20,000 and 30,000 particles seem to present a good trade-off between accuracy in tracking state variables and computational time.

4.3 Small Measurement Errors

In this section, the log-likelihood is simulated under the baseline model and a new set of observable variables. The objective of this section is to evaluate both filters when the signal-to-noise ratio on data is high. The new observable variables are built with the same simulated states, but less volatile measurement errors. In previous exercises, measurement errors have the same standard-deviation as the ergodic values for each variable; in this section, standard-deviation of

measurement errors are reduced to 10% of the ergodic standard-deviation of each variable.

Table 7 summarizes information on the effective number of particles, Monte Carlo variation of each filter and the speedup across filters. The average loss between effective and actual number of particles increases in this framework to 94% and 90% in the Bootstrap Filter and in the Auxiliary Particle Filter, respectively. Despite smaller losses on average, the Auxiliary Particle Filter seems to miss correct path of states. Notice, first, how both filters eventually collapse, with a degenerate effective particle characterizing the state of distribution. Second, in terms of Monte Carlo variation, while the Bootstrap Filter converges as the number of particles in the system increases, the Auxiliary Particle Filter presents larger and increasing losses as the size of the swarm grows. This result is a consequence of the joint distribution of measurement errors, which is too narrow in the model with five observables. In alternative tests, available upon request, the Auxiliary Particle Filter presents the proper and expected convergence when only output is used as an observable variable.

Table 7: Effective Particles – Small Measurement Errors

Samples		$\widehat{N^{eff}}$	$\min\left(\widehat{N^{eff}}\right)$	$\frac{Std(L)}{\mu(L)}(\%)$	Seconds per Sample	Speedup
1000	B.Filter	53.35	1	0.233	1.2263	0.7733
	Aux.Filter	93.10	1	0.956	1.6547	
5000	B.Filter	259.65	1	0.095	1.3299	0.8131
	Aux.Filter	466.89	1	0.991	1.9444	
10000	B.Filter	516.98	1	0.066	1.7608	0.8742
	Aux.Filter	935.64	1	1.000	2.7857	
20000	B.Filter	1031.0	1	0.046	2.3083	0.9642
	Aux.Filter	1874.8	1	1.076	4.0471	
30000	B.Filter	1545.1	1	0.037	2.8816	1.0403
	Aux.Filter	2814.8	1	1.113	5.4612	
40000	B.Filter	2059.1	1	0.032	3.6425	1.0482
	Aux.Filter	3753.8	1	1.194	6.9607	
50000	B.Filter	2573.2	1.01	0.028	4.1708	1.1000
	Aux.Filter	4693.7	1	1.196	8.3686	

In terms of speedup, the Auxiliary Particle Filter, for the first time, presents a similar performance compared to the Bootstrap Filter. There is a 23% gain in terms of seconds per effective particles for the smallest particle swarm and the two filters have similar performance at 30,000 particles. At the largest particle swarm, the Auxiliary Particle Filter presents a loss of only 10%, value unmatched by any size of the particle swarm in the baseline exercise presented in table 3.

Figure 6 shows the MAE for the model with small measurement errors. Results associating smaller gains in terms of MAE to a swarm with more than 20,000 particles are still robust, as expected. When compared to the outcome of figure 3, the increase in the signal-to-noise ratio is reflected by the overall increased precision of the predicted state variables. This result holds irrespective of the particle filter used, which is a little surprising given the poor performance of the Auxiliary Particle Filter in terms of Monte Carlo variation.

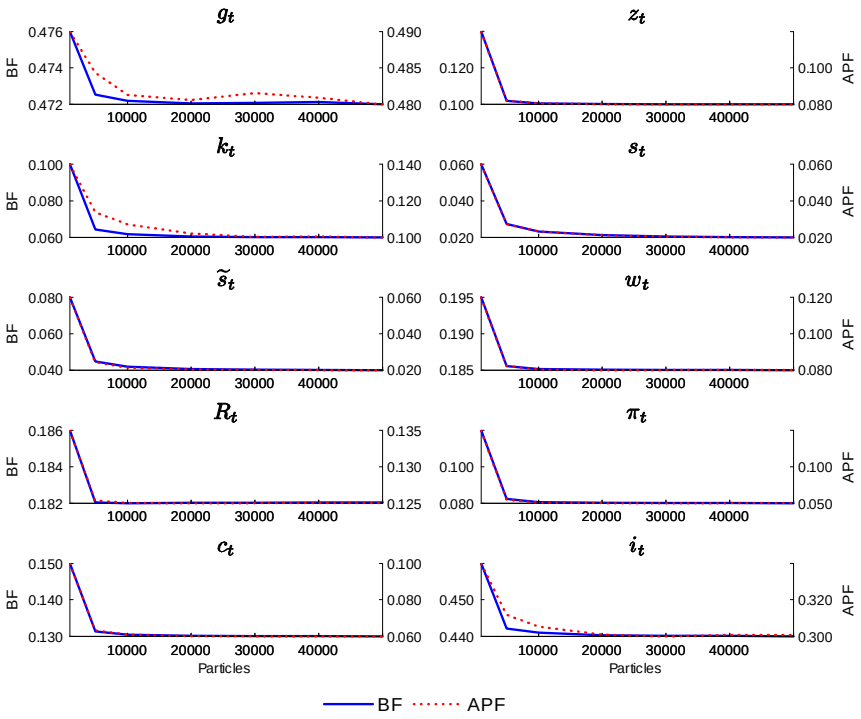

Figure 6: Mean Absolute Error (%) –State Variables – Small Measurement Errors

5 Conclusion

This note presented an evaluation of the number of particles necessary to properly approximate the likelihood of a prototypical DSGE model. It also compared the performance of two particle filters usually seen in the literature: the Bootstrap Filter and the Auxiliary Particle Filter. Results based on the accuracy of the simulated likelihood and the tracking of state variables show that a particle filter with a swarm of 20,000 and 30,000 particles seem to provide good performance, irrespective of the size of the model.

In the comparison of the two filters, the Auxiliary Particle Filter provided good performance when the number of observables in the model is large and the size of the particle swarm is relatively small. This is due to the fact that, under the baseline setup of the exercise, the Auxiliary Particle Filter provides a better characterization of the tails of the state vector distribution. On the other hand, the Bootstrap Filter is significantly less demanding in terms of computational resources. Thus, it might be possible, depending on the setup of the problem, to obtain better performance in the Bootstrap Filter with a large number of particles in the swarm, instead of using the Auxiliary Particle Filter.

The poor performance of the Auxiliary Particle Filter under alternative settings of the prob-

lem, combined with the high computational requirements, demands a deeper investigation in terms of alternative proposal distributions for states in the prediction step of the filter. The use of current information to draw states generate most of the time generate gains in the performance of particle filters. It should not be different in the problem proposed in this paper. Results, however, show that there might room to improve performance if different proposals are adopted to draw states. Also, the final choice between the Auxiliary Particle Filter and the Bootstrap Filter is still conditional to the computational resources available and the type of problem being solved. This note tried to provide some guidelines in the choice between these alternatives.

References

[1] S. A. L. Alves. *Essays on trend inflation: Optimal policy, estimation bias, and unemployment volatility*. PhD thesis, University of California, Santa Cruz, Oct. 2011.

[2] S. An. Bayesian estimation of DSGE models: Lessons from second-order approximations. Bank for International Settlements - BIS, June 2008.

[3] S. An and F. Schorfheide. Bayesian analysis of DSGE models. *Econometric Reviews*, 26(2-4):113–172, 2007.

[4] M. M. Andreasen. Non-linear DSGE models and the optimized central difference particle filter. *Journal of Economic Dynamics and Control*, 35(10):1671–1695, Oct. 2011.

[5] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.

[6] S. B. Aruoba, J. Fernández-Villaverde, and J. F. Rubio-Ramírez. Comparing solution methods for dynamic equilibrium economies. *Journal of Economic Dynamics and Control*, 30(12):2477–2508, 2006.

[7] Y. Boers. On the number of samples to be drawn in particle filtering. In *IEE Colloquium on Target Tracking: Algorithms and Applications (Ref. No 1999/090, 1999/215)*, pages 5/1–5/6. IET, 1999.

[8] M. Bolic, P. M. Djuric, and S. Hong. Resampling algorithms and architectures for distributed particle filters. *IEEE Transactions on Signal Processing*, 53(7):2442—2450, July 2005.

[9] L. J. Christiano, M. Eichenbaum, and C. L. Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45, Feb. 2005.

[10] D. Crisan and A. Doucet. A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on Signal Processing*, 50(3):736–746, Mar. 2002.

[11] J. Fernández-Villaverde, P. Guerrón-Quintana, and J. F. Rubio-Ramírez. Fortune or virtue: Time-Variant volatilities versus parameter drifting in U.S. data. NBER Working Papers 15928, National Bureau of Economic Research, Apr. 2010.

[12] J. Fernández-Villaverde and J. F. Rubio-Ramírez. Estimating nonlinear dynamic equilibrium economies: a likelihood approach. Working Paper 2004-1, Federal Reserve Bank of Atlanta, 2004.

[13] J. Fernández-Villaverde and J. F. Rubio-Ramírez. Estimating dynamic equilibrium economies: linear versus nonlinear likelihood. *Journal of Applied Econometrics*, 20(7):891–910, 2005.

[14] J. Fernández-Villaverde and J. F. Rubio-Ramírez. Estimating macroeconomic models: A likelihood approach. *Review of Economic Studies*, 74(4):1059–1087, 2007.

[15] J. Fernández-Villaverde, J. F. Rubio-Ramírez, and M. S. Santos. Convergence properties of the likelihood of computed dynamic models. *Econometrica*, 74(1):93–119, 2006.

[16] T. Flury and N. Shephard. Bayesian inference based only on simulated likelihood: Particle filter analysis of dynamic economic models. *Econometric Theory*, 27(05):933–956, 2011.

[17] J. Geweke. Bayesian inference in econometric models using monte carlo integration. *Econometrica*, 57(6):1317–39, 1989.

[18] N. J. Gordon, D. J. Salmond, and A. F. Smith. Novel approach to nonlinear/non-Gaussian bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F*, 140(2):107–113, Apr. 1993.

[19] J. D. Hamilton. *Time series analysis*. Princeton University Press, 1994.

[20] K. L. Judd. *Numerical methods in economics*. MIT Press, 1998.

[21] O. Kamenik. DSGE models with Dynare++. A tutorial. Technical report, Jan. 2011.

[22] H. R. Künsch. Recursive monte carlo filters: Algorithms and theoretical analysis. *The Annals of Statistics*, 33(5):1983–2021, Oct. 2005.

[23] J. S. Liu and R. Chen. Blind deconvolution via sequential imputations. *Journal of the American Statistical Association*, 90(430):567–576, June 1995.

[24] M. J. Lombardi and S. Sgherri. (Un)naturally low? sequential monte carlo tracking of the US natural interest rate. Working Paper Series 794, European Central Bank, 2007.

[25] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599, June 1999.

[26] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer, 2. ed. edition, 2004.

[27] S. Schmitt-Grohé and M. Uribe. Optimal operational monetary policy in the Christiano-Eichenbaum-Evans model of the U.S. business cycle. NBER Working Papers 10724, National Bureau of Economic Research, Aug. 2004.

[28] F. Smets and R. Wouters. An estimated dynamic stochastic general equilibrium model of the euro area. *Journal of the European Economic Association*, 1(5):1123–1175, 2003.

[29] M. Smith. Estimating nonlinear economic models using surrogate transitions. In *Workshop on Methods and Applications for Dynamic Stochastic General Equilibrium Models*. Federal Reserve Bank of Philadelphia, Oct. 2011.

[30] T. Yun. Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics*, 37(2-3):345–370, 1996.

Banco Central do Brasil

Trabalhos para Discussão

Os Trabalhos para Discussão do Banco Central do Brasil estão disponíveis para download no website http://www.bcb.gov.br/?TRABDISCLISTA

Working Paper Series

The Working Paper Series of the Central Bank of Brazil are available for download at http://www.bcb.gov.br/?WORKINGPAPERS

245	**Pesquisa Trimestral de Condições de Crédito no Brasil** *Clodoaldo Aparecido Annibal e Sérgio Mikio Koyama*	Jun/2011
246	**Impacto do Sistema Cooperativo de Crédito na Eficiência do Sistema Financeiro Nacional** *Michel Alexandre da Silva*	Ago/2011
247	**Forecasting the Yield Curve for the Euro Region** *Benjamim M. Tabak, Daniel O. Cajueiro and Alexandre B. Sollaci*	Aug/2011
248	**Financial regulation and transparency of information: first steps on new land** *Helder Ferreira de Mendonça, Délio José Cordeiro Galvão and Renato Falci Villela Loures*	Aug/2011
249	**Directed clustering coefficient as a measure of systemic risk in complex banking networks** *B. M. Tabak, M. Takami, J. M. C. Rocha and D. O. Cajueiro*	Aug/2011
250	**Recolhimentos Compulsórios e o Crédito Bancário Brasileiro** *Paulo Evandro Dawid e Tony Takeda*	Ago/2011
251	**Um Exame sobre como os Bancos Ajustam seu Índice de Basileia no Brasil** *Leonardo S. Alencar*	Ago/2011
252	**Comparação da Eficiência de Custo para BRICs e América Latina** *Lycia M. G. Araujo, Guilherme M. R. Gomes, Solange M. Guerra e Benjamin M. Tabak*	Set/2011
253	**Bank Efficiency and Default in Brazil: causality tests** *Benjamin M. Tabak, Giovana L. Craveiro and Daniel O. Cajueiro*	Oct/2011
254	**Macroprudential Regulation and the Monetary Transmission Mechanism** *Pierre-Richard Agénor and Luiz A. Pereira da Silva*	Nov/2011
255	**An Empirical Analysis of the External Finance Premium of Public Non-Financial Corporations in Brazil** *Fernando N. de Oliveira and Alberto Ronchi Neto*	Nov/2011
256	**The Self-insurance Role of International Reserves and the 2008-2010 Crisis** *Antonio Francisco A. Silva Jr*	Nov/2011

257	**Cooperativas de Crédito: taxas de juros praticadas e fatores de viabilidade** *Clodoaldo Aparecido Annibal e Sérgio Mikio Koyama*	Nov/2011
258	**Bancos Oficiais e Crédito Direcionado – O que diferencia o mercado de crédito brasileiro?** *Eduardo Luis Lundberg*	Nov/2011
259	**The impact of monetary policy on the exchange rate: puzzling evidence from three emerging economies** *Emanuel Kohlscheen*	Nov/2011
260	**Credit Default and Business Cycles: an empirical investigation of Brazilian retail loans** *Arnildo da Silva Correa, Jaqueline Terra Moura Marins, Myrian Beatriz Eiras das Neves and Antonio Carlos Magalhães da Silva*	Nov/2011
261	**The relationship between banking market competition and risk-taking: do size and capitalization matter?** *Benjamin M. Tabak, Dimas M. Fazio and Daniel O. Cajueiro*	Nov/2011
262	**The Accuracy of Perturbation Methods to Solve Small Open Economy Models** *Angelo M. Fasolo*	Nov/2011
263	**The Adverse Selection Cost Component of the Spread of Brazilian Stocks** *Gustavo Silva Araújo, Claudio Henrique da Silveira Barbedo and José Valentim Machado Vicente*	Nov/2011
264	**Uma Breve Análise de Medidas Alternativas à Mediana na Pesquisa de Expectativas de Inflação do Banco Central do Brasil** *Fabia A. de Carvalho*	Jan/2012
265	**O Impacto da Comunicação do Banco Central do Brasil sobre o Mercado Financeiro** *Marcio Janot e Daniel El-Jaick de Souza Mota*	Jan/2012
266	**Are Core Inflation Directional Forecasts Informative?** *Tito Nícias Teixeira da Silva Filho*	Jan/2012
267	**Sudden Floods, Macroprudention Regulation and Stability in an Open Economy** *P.-R. Agénor, K. Alper and L. Pereira da Silva*	Feb/2012
268	**Optimal Capital Flow Taxes in Latin America** *João Barata Ribeiro Blanco Barroso*	Mar/2012
269	**Estimating Relative Risk Aversion, Risk-Neutral and Real-World Densities using Brazilian Real Currency Options** *José Renato Haas Ornelas, José Santiago Fajardo Barbachan and Aquiles Rocha de Farias*	Mar/2012
270	**Pricing-to-market by Brazilian Exporters: a panel cointegration approach** *João Barata Ribeiro Blanco Barroso*	Mar/2012
271	**Optimal Policy When the Inflation Target is not Optimal** *Sergio A. Lago Alves*	Mar/2012

272	**Determinantes da Estrutura de Capital das Empresas Brasileiras: uma abordagem em regressão quantílica** *Guilherme Resende Oliveira, Benjamin Miranda Tabak, José Guilherme de Lara Resende e Daniel Oliveira Cajueiro*	Mar/2012
273	**Order Flow and the Real: Indirect Evidence of the Effectiveness of Sterilized Interventions** *Emanuel Kohlscheen*	Apr/2012
274	**Monetary Policy, Asset Prices and Adaptive Learning** *Vicente da Gama Machado*	Apr/2012
275	**A geographically weighted approach in measuring efficiency in panel data: the case of US saving banks** *Benjamin M. Tabak, Rogério B. Miranda and Dimas M. Fazio*	Apr/2012
276	**A Sticky-Dispersed Information Phillips Curve: a model with partial and delayed information** *Marta Areosa, Waldyr Areosa and Vinicius Carrasco*	Apr/2012
277	**Trend Inflation and the Unemployment Volatility Puzzle** *Sergio A. Lago Alves*	May/2012
278	**Liquidez do Sistema e Administração das Operações de Mercado Aberto** *Antonio Francisco de A. da Silva Jr.*	Maio/2012
279	**Going Deeper Into the Link Between the Labour Market and Inflation** *Tito Nícias Teixeira da Silva Filho*	May/2012
280	**Educação Financeira para um Brasil Sustentável** **Evidências da necessidade de atuação do Banco Central do Brasil em educação financeira para o cumprimento de sua missão** *Fabio de Almeida Lopes Araújo e Marcos Aguerri Pimenta de Souza*	Jun/2012