

Bayesian Estimation of State Space Models Using Moment Conditions*

A. Ronald Gallant
Penn State University

Raffaella Giacomini
University College London

Giuseppe Ragusa
Luiss University

First draft: January 31, 2015

This draft: August 28, 2016

*Address correspondence to A. Ronald Gallant, P.O. Box 659, Chapel Hill NC 27514, USA, phone 919-428-1130; email aronldg@gmail.com.

We thank Gary Chamberlain, Whitney Newey, Frank Schorfheide, Neil Shephard and seminar participants at Harvard/MIT, Cambridge, Carlos III, Yale, Duke, Northern Illinois, UCSC, Penn State, Michigan State, Cleveland Fed, ULB, and Illinois for useful comments and suggestions. Raffaella Giacomini gratefully acknowledges financial support from the ESRC through the Centre for Microdata Methods and Practice grant RES-589-28-0001.

© 2015 A. Ronald Gallant, Raffaella Giacomini, and Giuseppe Ragusa.

Abstract

We consider Bayesian estimation of state space models when the measurement density is not available but estimating equations for the parameters of the measurement density are available from moment conditions. The most common applications are partial equilibrium models involving moment conditions that depend on dynamic latent variables (e.g., time-varying parameters, stochastic volatility) and dynamic general equilibrium models when moment equations from the first order conditions are available but computing an accurate approximation to the measurement density is difficult.

Keywords and Phrases: State Space Models, Bayesian Estimation, Moment Equations, Structural Models, DSGE Models, Particle Filter

JEL Classification: C32, C36, E27

1 Introduction

We propose a method for conducting Bayesian inference regarding the parameters of a non-linear structural model that has dynamic latent variables. By latent variables we mean all endogenous and exogenous variables in the model that are not observed.

The general approach to dealing with dynamic latent variables in econometrics is to resort to filtering techniques (e.g., the particle filter), which, in connection with Markov Chain Monte Carlo (MCMC) methods, deliver estimates of the structural parameters (see Andrieu, Doucet, and Holenstein, 2010). To implement a particle filter one needs to be able to: 1) draw from the transition density of the latent variables, which specifies the distribution of the latent variables conditional on their past history; and 2) evaluate the measurement density, which specifies the distribution of the observable variables conditional on the latent variables.

In this paper, we maintain the assumption that one can draw from the transition density of the latent variables but we assume that a measurement density is not available and/or it is difficult to approximate numerically. What is available is instead a set of moment conditions that provide estimating equations for the parameters of the measurement density. The most common applications in econometrics where this situation arises are 1) partial equilibrium models that involve moment conditions depending on dynamic latent variables (e.g., time-varying parameters, stochastic volatility); and 2) dynamic general equilibrium structural models when moment equations from the first order conditions are available but computing an accurate approximation to the measurement density is difficult. There are currently no econometric methods that apply to the first class of models, and for the second class of models our method can be considered as an alternative to existing approaches that does not rely on approximations or numerical solutions of the model.

The method of moments has a powerful appeal in economic research and researchers are increasingly keen to use prior information as a means to deal with data limitations. The method we propose here has potential to become a useful tool in applied economic research, because - as argued by Cochrane (2005) - most researchers find evidence based on method of moments more persuasive than evidence based on fully specified likelihoods. Our

contribution is to show that combining method of moments and priors is viable theoretically and practically in economic models where the presence of dynamic latent variables makes it impossible to apply standard GMM estimation.

In fact, if one considers calibration to be Bayesian method of moments with extremely strong priors, then most of the science that matters in our daily lives uses Bayesian method of moments. In particular, climate models and macro models. The main exception is health, but this is mostly due to government regulation. Also, the exceptions one finds in macro are mostly due to the pressure of central banks. Our view is that if statistics is to become relevant to major policy decisions, then something along the lines of what we propose has to become viable.

We illustrate the usefulness of our method by applying it to the problem of estimating the latent endowment process in a Lucas (1978) economy given only knowledge of the agent’s first order conditions and of the transition density of the latent process. The process we extract differs markedly from measured consumption and suggests the presence of stochastic volatility and jumps.

The central idea of the paper is to show that the moment conditions can be used to construct a “GMM representation” of the measurement density that one can substitute for the measurement density as an input into an otherwise standard filtering MCMC algorithm.

To illustrate, suppose we have a set of M moment conditions

$$\mathcal{E}[g(y_{t+1}, x_{t+1}, \theta)] = 0^1$$

implied by a structural model. We observe a realization $y = \{y_1, \dots, y_T\}$ from the stochastic process $\{\dots, y_{t-1}, y_t, y_{t+1}, \dots\}$ but we do not observe $\{\dots, x_{t-1}, x_t, x_{t+1}, \dots\}$ which is thus the latent process. What we know about the latent process is a parametric specification for its transition density. The objective is to obtain the posterior distribution of the structural parameter θ (comprised of the parameters of both the moment conditions and the transition density) and the posterior distribution of the latent process. Formally, the posterior is given

¹Expectation for moment conditions is determined by context. If θ is regarded as exogenous and a likelihood $p(x, y | \theta)$ is well defined, then the meaning is $\iint g(y_{t+1}, x_{t+1}, \theta) p(x, y | \theta) dy dx = 0$. If θ is regarded as endogenous, then the meaning is $\iiint g(y_{t+1}, x_{t+1}, \theta) p(x, y, \theta) dy dx d\theta = 0$. Equation (43) of the application in Section 6 is an instance of the latter case. The examples in Section 5 are instances of the former.

by

$$p^o(\theta, x|y) \propto p^o(y|x, \theta)p^o(x|\theta)p^o(\theta)$$

where the measurement density $p^o(y|x, \theta)$ is unknown aside from the restrictions implicitly imposed by the moment conditions, the joint density of the latent variables $p^o(x|\theta)$ is pinned down by the transition density, and the prior $p^o(\theta)$ of the parameters is specified by the researcher. The contribution of this paper is twofold. We first show that the moment conditions induce a probability structure that allows us to replace the unknown transition density $p^o(y|x, \theta)$ with a known density $p^*(y|x, \theta)$. We then propose a numerical algorithm that uses the particle filter and a Metropolis algorithm to draw from the posterior $p^*(\theta, x|y) \propto p^*(y|x, \theta)p^o(x|\theta)p^o(\theta)$.

Regarding the first contribution, we build on and extend the results of Gallant and Hong (2007) and Gallant (2016a, 2016b, 2016c) to an environment with dynamic latent variables. The key insight is to show how to replace the probability space over $(\mathcal{Y} \times \mathcal{X} \times \Theta, \mathcal{C}^o, P^o)$ implied by the structural model and a prior for θ (where $\mathcal{Y} \times \mathcal{X}$ is the support of the observable and latent variables, Θ is the support of θ , and \mathcal{C}^o is the collection of Borel subsets of $\mathcal{Y} \times \mathcal{X} \times \Theta$) by an alternative probability space $(\mathcal{Y} \times \mathcal{X} \times \Theta, \mathcal{C}^*, P^*)$. The alternative probability space is such that \mathcal{C}^* is a subset of \mathcal{C}^o and the density of P^* is the same as P^o except that the measurement density is replaced by a density function evaluated at the sample moment conditions g_T (scaled to have variance equal to the identity matrix, i.e., $p^*(y|x, \theta) = \psi([\Sigma(y, x, \theta)]^{-1/2} g_T(y, x, \theta))$). We call this density function the ‘‘GMM representation’’ of the measurement density. Because we are concerned with subjective Bayesian inference, we assume that the density function ψ is specified by the user.² In practice, we suggest using the standard normal density, which is motivated by the asymptotic normality of the sample moments under the standard regularity assumptions. The key insight that allows us to substitute the unknown measurement density with its GMM representation is the fact that both probability measures assign the same probability to sets in \mathcal{C}^* . Naturally, because \mathcal{C}^* is a subset of \mathcal{C}^o , some information is lost. Intuitively this is similar to the information loss that occurs when one divides the range of a continuous variable into intervals and uses a

²While this article was in press, methods for determining Ψ from primitives were proposed in Gallant (2016d).

discrete distribution to assign probability to each interval. Both the continuous and discrete distributions assign the same probability to each interval but the discrete distribution cannot assign probability to subintervals. How much information is lost depends on how well one chooses moment conditions. An in-depth investigation of the effects of moment choice on inference is beyond the scope of this paper, but we provide some advice on choice strategy for some key economic applications. In many instances, as in the application of Section 6, discussion of the choice of moments is moot because the economics of the situation dictate the choice.

In the state-space literature to which we contribute, (cf. Flury and Shephard (2011), Fernandez-Villaverde and Rubio-Ramirez (2006)) the assumption that one can draw from the transition density is standard. Our contribution is to be able to perform Bayesian inference without knowledge of the measurement density.

The importance of the first contribution is easy to overlook. What it does is establish the methodology as exact within the Bayesian paradigm given the information that the researcher chooses to use. Leaving aside specification error, inaccurate algorithms, etc. that plague all statistical methods, we are proposing exact Bayesian methods, not approximate Bayesian methods.

Regarding our second contribution, which builds on ideas from Beaumont (2003), Andrieu and Roberts (2009), Andrieu, Doucet, and Holenstein (2010), and Flury and Shephard (2011), the computational strategy we propose consists of two steps: a conditional particle filter step that draws x given y , θ , and the previously drawn x and a Metropolis step that draws θ given y , x , and the previously drawn θ . The validity of the algorithm follows from the results of Andrieu, Doucet, and Holenstein (2010) as it can be thought of as an adaptation of their particle Gibbs sampler when one has to resort to the GMM representation of the measurement density. The application of the algorithm results in an MCMC chain in (θ, x) and thus parameter estimates, standard deviations, and other characterizations of the posterior distribution can be computed from this chain in the standard way (Gamerman and Lopes, 2006).

The main attraction of the method we propose is that one does not have to solve the structural model. For partial equilibrium models this is crucial because, in general, there do

not exist practicable alternatives.

We also expect that an important application for our results will be statistical inference regarding general equilibrium models in macroeconomic applications such as dynamic stochastic general equilibrium models (DSGE). For analytically intractable DSGE models there are alternatives to what we propose that rely on being able to solve the model numerically. For instance, one can use perturbation methods to approximate the model, use the approximation to obtain an analytical expression for the measurement density, and then use some method of numerical integration such as particle filtering to eliminate the latent variables along the lines proposed by Fernandez-Villaverde and Rubio-Ramirez (2006) and Flury and Shephard (2011). Alternatively, one can solve the model only to the point of being able to simulate it and then use the methods proposed by either Gallant and McCulloch (2009), who use an SNP (Gallant and Nychka, 1987) representation of the measurement density, or Gallant and Tauchen (2015), who use an EMM (Gallant and Tauchen, 1996) representation of the measurement density.

In the case of DSGE models, the main reason one might want to consider our alternative to the existing procedures is that one has misgivings about the quality of the numerical methods one has used to solve the structural model. For instance, perturbation methods such as linearization cause loss of information: they typically require dealing with singularity issues and with possible multiplicity of solutions (indeterminacy). Moreover, lower order expansions can lose important features of a model such as stochastic volatility (Bloom, 2009; Benigno, Benigno, Nisticó, 2012). A secondary reason is to avoid singularities in the measurement equation that can arise when using a likelihood based approach with particle filtering; see, e.g., Subsection 5.2.

2 Assumptions and Implications

ASSUMPTION 1 We require the existence of (but not complete knowledge of) a dynamic structural model that has parameter $\theta \in \Theta$. We observe $y = (y_1, y_2, \dots, y_T) \in \mathcal{Y}$, a subset of the endogenous and exogenous variables in the model. We do not observe the variables in the model that remain: $x = (x_1, x_2, \dots, x_T) \in \mathcal{X}$. These are the latent variables. Partial histories are denoted $y_{1:t} = (y_1, y_2, \dots, y_t)$ and $x_{1:t} = (x_1, x_2, \dots, x_t)$. The variables y_t and

x_t are vectors, as is θ . ■

ASSUMPTION 2 The set $\mathcal{Y} \times \mathcal{X} \times \Theta$ is a Borel subset of $\mathbb{R}^{\dim(y)+\dim(x)+\dim(\theta)}$. Denote the probability measure over the Borel subsets of $\mathcal{Y} \times \mathcal{X} \times \Theta$ implied by the structural model and the prior for θ by P° . P° has joint density density $p^\circ(y, x, \theta)$ that factors as

$$p^\circ(y, x, \theta) = p^\circ(y | x, \theta)p^\circ(x | \theta)p^\circ(\theta), \quad (1)$$

where $p^\circ(x | \theta) = \left[\prod_{t=2}^T p^\circ(x_t | x_{t-1}, \theta) \right] p^\circ(x_1 | \theta)$. We do not assume knowledge of the measurement density $p^\circ(y | x, \theta)$. We assume that the process $\{x_t\}_{t=-\infty}^\infty$ is ergodic, that we can draw from its transition density $p^\circ(x_t | x_{t-1}, \theta)$, and that we can evaluate the prior $p^\circ(\theta)$. ■

If $p^\circ(y, x, \theta)$ exists, the factorization (1) is always available. Interest in this article is focused on the case where one knows $p^\circ(x | \theta)$ and $p^\circ(\theta)$ but not $p^\circ(y | x, \theta)$ and therefore not $p^\circ(y, x, \theta)$. Note that we can draw from the stationary density $p^\circ(x_1 | \theta)$ of the process $\{x_t\}_{t=-\infty}^\infty$ by drawing from the transition density with an arbitrary start x_0 and waiting for transients to die out.

Examples of latent variables that satisfy Assumption 2 and are routinely used in economic models are time-varying parameters, structural shocks, state-dependent parameters, and state-dependent factors. If necessary to accommodate state dependence, e.g., Markov switching, one can modify the functional form of the transition density provided that ergodicity is retained. Therefore, the transition density could, e.g., be of the form $p^\circ(x_{t+1} | x_{1:t}, y_{1:t}, \theta)$. However, in this case, one must be able to evaluate the transition density and the formulas for the weights in the particle filters we describe later become more complicated. On this see Gallant, Hong, and Khwaja (2015).

When working with DSGE models one is used to thinking in terms of observables and states. That is not the dichotomy we have in mind here. Our division is into what is observed and what is not observed. Thus, what we term latent variables can include unobserved states, unobserved exogenous variables, and unobserved endogenous variables. The practical limit on what is permitted is determined by Assumption 2.

ASSUMPTION 3 The structural model implies a set of moment conditions of the form

$$\mathcal{E}[g(y_t, x_t, \theta)] = \iiint g(y_t, x_t, \theta) p^\circ(y, x, \theta) dy dx d\theta = 0, \quad (2)$$

where $g : (y_t, x_t, \theta) \mapsto \mathbb{R}^M$. ■

REMARK 1 Usually, but not always, the derivation of (2) proceeds as follows. The parameter θ is not regarded as endogenous so that the structural model has a well defined likelihood $p^o(y, x | \theta)$ that we shall assume is Markovian for convenience. The structural model implies a conditional moment condition of the form

$$\iint h(y_{t+1}, x_{t+1}, \theta) p^o(y_{t+1}, x_{t+1} | y_t, x_t, \theta) dy_{t+1} dx_{t+1} = 0$$

that holds for every θ in the model's parameter space, where $h(y_{t+1}, x_{t+1}, \theta)$ is vector valued. Let $A(y_t, x_t)$ be matrix valued. Then, because $A(y_t, x_t)$ is orthogonal to $h(y_{t+1}, x_{t+1}, \theta)$,

$$\mathcal{E}[A(y_t, x_t)h(y_{t+1}, x_{t+1}, \theta)] = \iint A(y_t, x_t)h(y_{t+1}, x_{t+1}, \theta) p^o(y, x | \theta) dy dx = 0 \quad (3)$$

holds for every θ . Because (3) holds for every θ , (3) implies (2) for $g(y_t, x_t, \theta) = A(y_{t-1}, x_{t-1})h(y_t, x_t, \theta)$. □

Sample moments corresponding to (2) are

$$g_T(y, x, \theta) = \frac{1}{\sqrt{T}} \sum_{t=1}^T g(y_t, x_t, \theta) \quad (4)$$

with weighting matrix

$$\Sigma(y, x, \theta) = \frac{1}{T} \sum_{t=1}^T [\tilde{g}(y_t, x_t, \theta)][\tilde{g}(y_t, x_t, \theta)]' \quad (5)$$

$$\tilde{g}(y_t, x_t, \theta) = g(y_t, x_t, \theta) - \frac{1}{\sqrt{T}} g_T(y, x, \theta) \quad (6)$$

If the moment conditions are serially correlated one will have to substitute a heteroskedastic autoregressive consistent (HAC) weighting matrix (Andrews, 1991) for that shown as (5). Whether $\Sigma(y, x, \theta)$ is cross-sectional as in (5) or HAC, the residuals used to compute it should be of the form shown as (6).

ASSUMPTION 4 The sample moments normalized by the weighting matrix

$$Z = [\Sigma(y, x, \theta)]^{-1/2} g_T(y, x, \theta) \quad (7)$$

have a known distribution Ψ with density $\psi(z)$. When Z is computed from a partial history we write $Z_{1:t}$. Note that $Z : (y, x, \theta) \mapsto \mathbb{R}^M$ and $Z_{1:t} : (y_{1:t}, x_{1:t}, \theta) \mapsto \mathbb{R}^M$. ■

We are concerned with Bayesian subjective inference so that one asserts that Z follows Ψ rather than assumes although one should take care that the assertion is not contradicted by the primitives of the structural model. A standard choice for Ψ in applications is the standard normal distribution Φ with density $\phi(z) = (2\pi)^{M/2}e^{-\frac{1}{2}z'z}$. Regularity conditions such that (7) is asymptotically normal are in Hansen (1982), Gallant and White (1987), and elsewhere. When these regularity conditions are in force, asserting that Z is standard normal becomes more palatable, as does an assertion that Z has the multivariate Student- T should one wish to allow fatter tails to acknowledge a modest sample size. There are risks in following this approach for which see Gallant (2016b, 2016d). An asymptotic equicontinuity condition under which asymptotic normality of Ψ is necessary and sufficient for uniform (in x , y , and θ) convergence of the density ψ of Ψ is given by Theorem 3 of Sweeting (1986). If Sweeting's asymptotic equicontinuity condition can be verified, then this risk is considerably mitigated. A compact Θ , domination and smoothness conditions on $g(y_t, x_t, \theta)$, and bounds on the eigenvalues of the weighting matrix $\Sigma(y, x, \theta)$ usually figure among the regularity conditions that insure asymptotic normality. As seen later, these assumptions lead to a presumption that particle filter weights will be bounded and that the MCMC chain we propose will mix.

Our final assumption is critical. It provides Z with some of the properties of a pivotal.

ASSUMPTION 5 Let $\mathcal{Z} = \{z \in \mathbb{R}^M : \psi(z) > 0\}$ and

$$C^{(\theta, z)} = \{(y, x) \in \mathcal{Y} \times \mathcal{X} : Z(y, x, \theta) = z\}. \quad (8)$$

We assume that $C^{(\theta, z)}$ is not empty for any $(\theta, z) \in \Theta \times \mathcal{Z}$. ■

As yet we have not encountered a practical application that violates Assumption 5. Sufficient in the case $\mathcal{Z} = \mathbb{R}^M$ is that each element of g_T is continuous with respect to at least one continuous element of (y, x) , is neither bounded from above nor below as that continuous variable varies over its support, and that the residuals used to compute the weighting matrix are centered as in (6).

Assumption 5 allows us to replace the probability space

$$(\mathcal{Y} \times \mathcal{X} \times \Theta, \mathcal{C}^o, P^o)$$

given by Assumption 2 with a probability space

$$(\mathcal{Y} \times \mathcal{X} \times \Theta, \mathcal{C}^*, P^*)$$

where P^* has density

$$p^*(y, x, \theta) = p^*(y | x, \theta) p^o(x | \theta) p^o(\theta) \quad (9)$$

$$p^*(y | x, \theta) = \psi[Z(y, x, \theta)] \quad (10)$$

and $\mathcal{C}^* \subset \mathcal{C}^o$ for the purpose of Bayesian inference (Gallant, 2016a).

\mathcal{C}^* is defined as follows.³ Let \mathcal{C} be the smallest σ -algebra that contains the preimages $C = \{(y, x, \theta) : Z(y, x, \theta) \in B \cap \mathcal{Z}\}$ where B ranges over the Borel subsets of \mathbb{R}^M . \mathcal{C}^* is the smallest σ -algebra that contains all sets in \mathcal{C} plus all sets of the form $R_B = (\mathbb{R}^{\dim(y)} \times B) \cap (\mathcal{Y} \times \mathcal{X} \times \Theta)$, where B is a Borel subset of $\mathbb{R}^{\dim(x)} \times \mathbb{R}^{\dim(\theta)}$. Gallant (2016a, Section 3) proves that $P^o(C) = P^*(C)$ for all $C \in \mathcal{C}^*$.

There is a term missing in (10); i.e., the right hand side should read $\text{adj}(y, x, \theta)\psi[Z(y, x, \theta)]$. The normalizing constant $\text{adj}(y, x, \theta)$ is analogous to a Jacobian term. Its construction is described in Gallant (2016b, 2016d). If $\text{adj}(y, x, \theta)$ does not depend on θ , then it can be disregarded when using MCMC as we propose. When it does depend on θ , its omission can be interpreted as using a data dependent prior. Even so, Gallant (2016b) argues on the basis of interpretability of results that it is preferable to omit the adjustment in applications; i.e., set $\text{adj}(y, x, \theta) = 1$. Comparison of lines 2 and 5 of his Table 1 suggests that in samples of reasonable size, neglecting the adjustment does not markedly affect results.

As mentioned earlier, $\mathcal{C}^* \subset \mathcal{C}^o$ implies a possible loss of information. The choice of moment conditions governs how much, if any, information is lost. If they are chosen such that $\mathcal{C}^* = \mathcal{C}^o$, then no information is lost. Saying that the measurement density has a ‘‘GMM representation’’ is motivated by the fact that when Ψ is the standard normal distribution we have

$$p^*(y_{1:T}|x_{1:T}, \theta) = (2\pi)^{\frac{M}{2}} \exp \left\{ -\frac{1}{2} g_T(y_{1:T}, x_{1:T}, \theta)' [\Sigma(y_{1:T}, x_{1:T}, \theta)]^{-1} g_T(y_{1:T}, x_{1:T}, \theta) \right\},$$

³The correspondence with Gallant’s (2016a, Section 3) notation is that y is his random variable X and (x, θ) is his random variable Λ .

which is a transformation of the GMM objective function we could use to estimate θ were $x_{1:T}$ to be observed. See, e.g., Chernozhukov and Hong (2003).

Bayesian inference relies on the joint posterior density

$$p^*(\theta, x_{1:T}|y_{1:T}) \propto \psi [Z(y_{1:T}, x_{1:T}, \theta)] \prod_{t=1}^T p^o(x_t|x_{t-1}, \theta)p^o(\theta). \quad (11)$$

Because $\psi [Z(y, x, \theta)]$ can only be evaluated point-wise, it is necessary to resort to the Monte Carlo method to conduct inference about $(x_{1:T}, \theta)$. To this end, we will employ the Particle MCMC of Andrieu, Doucet, and Holenstein (2010). The key idea is that for a fixed value of θ the particle filter provides an unbiased estimate of the marginal likelihood

$$p^*(y_{1:T}) = \int \psi [Z(y_{1:T}, x_{1:T}, \theta)] \prod_{t=1}^T p^o(x_t|x_{t-1}, \theta)p^o(\theta) dx_{1:T} d\theta.$$

A possibility is to use the unbiased estimate of $p(y_{1:T})$ within a Metropolis-Hastings algorithm to target the joint posterior in (11). This is the Particle Marginal Metropolis-Hastings (PMMH) algorithm (Algorithm 2 of Section 3).

Instead of relying on this algorithm, the chain that we propose for Bayesian estimation, described in more detail in Section 3, is a Particle Gibbs sampler that draws from $p(\theta, X_{1:T})$ by sampling iteratively from $p(\theta|y_{1:T}, x_{1:T})$ and $p(x_{1:T}|\theta, y_{1:T})$. Since the density $p(\theta|y_{1:T}, x_{1:T})$ is intractable, we use a Metropolis step to draw θ . Notice that draws of $x_{1:T}$ need to be conditional on previously drawn x otherwise the algorithm would not target the correct posterior distribution (see, Andrieu, Doucet, and Holenstein, 2010). Such conditional draws can be obtained by using a particle filter based on a conditional Sequentially Monte Carlo (SMC).

2.1 Choice of Moments

There are several received facts regarding method of moments: 1) The choice of moments affects results; 2) Algebraic manipulation of moment conditions can affect results; e.g., using one moment equation to substitute out a variable in another; 3) Optimal moments span the scores of the density of observables given parameters. These we accept in what follows.

Often the science dictates the moments. That is, referring to Remark 1, economics dictates moments $A(y_t, x_t)h(y_{t+1}, x_{t+1}, \theta_{(1)})$ that identify the parameters $\theta_{(1)}$ of scientific

interest and there is little flexibility other than the choice of instruments $A(y_t, x_t)$. What remains is to add additional moments to identify x and the parameters $\theta_{(2)}$ of the transition density $p(x_t | x_{t-1}, \theta_{(2)})$. The scores for x_{t-1} and $\theta_{(2)}$ are the best choice. The application in Section 6 is an illustration of this situation.

We now argue on the basis of heuristics and experience that the moment conditions for the Gibbs and Metropolis steps can be different. Obviously, if one has misgivings about splitting moments into two groups, then one can elect to follow our recommendations for each group and combine the two groups into one and use the one group in both the conditional particle filter and Metropolis steps. Splitting moments into two groups reduces computation time. If, say, one can divide ten moment conditions into two groups of five each, then computation time would more than halve.

We advise the reader that the notion of splitting moments does not appeal to the referees. Our recommendations below are based on the received facts above, intuition, and experience, not formal proof. Our experience is that the particle filter step and the Metropolis step are most affected by different subsets of the entire set of moments. Intuition suggests that this is because each is targeting different variables conditional on knowing those not targeted. Therefore, if one can deliberately enhance this separation and then divide the moments into two groups, computational costs can be reduced and the accuracy of posterior credibility intervals improved.

For the particle filter step the perfect moments would be those that spanned the scores of the conditional density of y given x and θ , were they known. For the Metropolis step the perfect moments would be those that spanned the scores of the density for y given θ , were they known. These choices are not practicable without an analytic expression for $p^o(y, x, \theta)$, in which case there is no point to using our proposals. However, there do seem to be some principles one can apply in selecting moments at the Metropolis step that we have discovered in our experimentation.

For the Metropolis step, one should try to identify as many parameters as possible from the observed data alone and try to make the latent variables depend as much as possible on quantities that can be computed from the observed data. If one is successful at this, then estimation results will be satisfactory, in our experience, but particle filter draws from

the conditional distribution of the latent variables will not mimic the true (but unobserved) trajectory of the latent variables very well. This can be corrected, in our experience, by choosing the moments used in the particle filter step so that observed variables depend on the latent variables as much as possible without regard for identification of parameters. I.e., the exact opposite of the goal for choosing moments for the Metropolis step. We illustrate these principles in the DSGE example of Subsection 5.2. Oddly enough, if our DSGE example is not misleading, a poor choice of moments at the particle filter step does not materially degrade the performance of the estimator for θ , as seen in Subsection 5.2.

Some parameters of a model, particularly a DSGE model, may not be identified even if the correct likelihood involving only observables were known. When this problem occurs, the unidentified parameters must be calibrated or restricted by tight priors. The DSGE example in Subsection 5.2 exhibits this problem and we deal with it by calibration.

3 Algorithms

In this section we present the particle Gibbs algorithm that we use in our applications. We also discuss the PMMH algorithm which, as said in the previous section, could also be used to sample from $p^*(\theta, x_{1:T}|y_{1:T})$, but in our experience does not work as well as the Gibbs method. We previously introduced the notation $y_{1:t} = (y_1, \dots, y_t)$, $x_{1:t} = (x_1, \dots, x_t)$, and $Z_{1:t}$. The densities $p^*(y_{1:t} | x_{1:t}, \theta)$ and $p^*(y_{1:t}, x_{1:t}, \theta)$ for partial histories are

$$p^*(y_{1:t}, x_{1:t}, \theta) = p^*(y_{1:t} | x_{1:t}, \theta)p^o(x_{1:t} | \theta)p^o(\theta) \quad (12)$$

$$p^*(y_{1:t} | x_{1:t}, \theta) = \psi[Z_{1:t}(y_{1:t}, x_{1:t}, \theta)]. \quad (13)$$

Three building blocks are required to implement the particle Gibbs algorithm. These are:

- A particle filter algorithm⁴
 - Input: θ .
 - Output: Draws $\{x_{1:T}^i\}_{i=1}^N$ from $p^*(x_{1:T} | y_{1:T}, \theta)$.
- A conditional particle filter step.

⁴Used for PMMH, Algorithm 2, or to generate counter-factuals given θ .

- Input: The previous draw $x_{1:T}^{(j-1)}$ and a draw $\theta^{(j)}$ from $p^*(\theta | y_{1:T}, x_{1:T}^{(j-1)})$.
- Output: A draw $x_{1:T}^{(j)}$ from $p^*(x_{1:T} | y_{1:T}, \theta^{(j)})$ that is conditional on $x_{1:T}^{(j-1)}$.
- A Metropolis-Hastings step.
 - Input: The previous draw $\theta^{(j)}$ and a draw $x_{1:T}^{(j)}$ from $p^*(x_{1:T} | y_{1:T}, \theta^{(j)})$.
 - Output: A draw $\theta^{(j+1)}$ from $p^*(\theta | y_{1:T}, x_{1:T}^{(j)})$ via a chain started at $\theta^{(j)}$.

While the particle filter, Algorithm 1, gives draws from the conditional distribution of $x_{1:T}$ given $y_{1:T}$ and θ that is memoryless with respect to previous draws of $x_{1:T}$, the conditional particle filter step of Algorithm 3 produces a draw from the conditional distribution of $x_{1:T}$ given $y_{1:T}$ and θ with memory of the previously drawn $x_{1:T}$. This dependence on the past trajectory of $x_{1:T}$ is necessary to guarantee that the resulting algorithm targets the right posterior. The last step is a basic Metropolis-Hasting algorithm that produces a draw from the conditional distribution of θ given $x_{1:T}$ and $y_{1:T}$.

Algorithm 1 about here

The particle filter is presented in Algorithm 1. It is a standard algorithm where the importance weights are constructed using $p^o(x_t | x_{t-1}; \theta)$ as proposal. More efficient algorithms that take into account current partial histories $y_{1:t}$ and/or future values of the latent variables are available. We do not consider these extensions of the particle filters since our main focus is to study the quality of inference when the measurement density has a GMM representation. However, it seems feasible to extend existing methods to obtain a more informative proposal such as the auxiliary particle filter of Pitt and Shephard (1999). The particle filter described in Algorithm 1 gives an unbiased estimator of $p^*(y_{1:T})$, viz.,

$$\hat{p}^*(y_{1:T} | \theta) = \prod_{t=1}^T \frac{1}{N} \sum_{i=1}^N w_t^i(\theta). \quad (14)$$

The PMMH sampler is an MCMC algorithm based on this unbiased estimate of the marginal distribution and it targets the full joint posterior distribution $p^*(\theta, x_{1:T} | y_{1:T})$. A description of this algorithm is given in Algorithm 2. An issue with the particle marginal Metropolis-Hastings samplers in high dimensions is that it often exhibits “sticky” behavior as an unusually “good” accepted $x_{1:T}$ is hard to displace. This motivates consideration of

a Gibbs-style algorithm where draws from $p^*(x_{1:T}|y_{1:T}, \theta)$ are always accepted. As reported by Chopin and Singh (2015), the particle Gibbs has favorable performance in practice, and this is also our experience.

Algorithm 2 about here

In an idealized Gibbs sampler, the posterior in (11) could be targeted by the following two steps (for $j = 1, \dots, R$)

1. Draw $\theta^{(j)}|x_{1:T}^{(j-1)} \sim p^*(\cdot|y_{1:T}, x_{1:T}^{(j-1)})$, and
2. Draw $x_{1:T}^{(j)}|\theta^{(j)} \sim p^*(\cdot|y_{1:T}, \theta^{(j)})$.

While in standard applications a draw from $p^*(\theta|y_{1:T}, x_{1:T})$ can sometime be carried out exactly, when the measurement density has the GMM representation one has to resort to Metropolis-Hastings. The second step could be accomplished by sampling from $p^*(x_{1:T}|y_{1:T}, \theta)$ using a standard particle filter like the one described in Algorithm 1. Unfortunately, this naive strategy would fail to target $p^*(\theta, x_{1:T}|\theta)$ (Andrieu, Doucet, and Holenstein, 2010). A remedy is to sample new particles conditional on a previous sampled trajectory. This can be accomplished by a slight modification of the particle filter described in Algorithm 1. The full algorithm that we implement to sample from $p^*(\theta, x_{1:T}|y_{1:T})$ is described in Algorithm 3.

Algorithm 3 about here

To implement the Metropolis-Hastings step we require a proposal density for θ . A proposal density is a transition density of the form $q(\theta|\theta^{(j-1)})$ such as a move-one-at-a-time random walk. In the examples of Section 5, we use the move-one-at-a-time random walk that uniformly selects an index k and then moves the element $\theta_k^{(j-1)}$ of $\theta^{(j-1)}$ to $\theta_k^{(j)}$ according to a normal with mean $\theta_k^{(j-1)}$ and variance σ_k , where σ_k is chosen by trial and error to achieve a rejection rate of about 50% in the Accept-Reject step of the algorithm that follows. In our implementation we also found useful to thin highly correlated chains by only writing every S th element of the chain; we refer to S as the “stride” when reporting results.

The good performance of the Particle Gibbs algorithm is backed up by theory. Chopin and Singh (2015) have recently shown that the particle Gibbs kernel is uniformly ergodic. The same paper shows that the particle Gibbs sampler can be made computationally more efficient by using a different way of conditioning on a past trajectory. As in the case of the particle filter, it is beyond the scope of the current paper to pursue such improvement, although we believe it could be useful to consider more sophisticated algorithms in future studies.

4 Theory

THEOREM 1 Under Assumptions 1 through 5 and mild additional regularity conditions, the particle Gibbs described in Algorithm 3 generates draws from $p^*(x_{1:T}, \theta | y_{1:T})$.

Proof

Regularity conditions sufficient to imply that particles are draws from the density $p^*(x | y, \theta)$ are in Andrieu, Doucet, and Holenstein (2010). They are mild, requiring that the weights at the importance sampling step be bounded and that multinomial resampling be used, which is the scheme used at the selection step.

In most applications regularity conditions such that the GMM estimator is consistent and asymptotically normal, were x and y observed, are in force. The regularity conditions used to prove consistency and asymptotic normality of GMM estimators typically include a compact parameter space, domination conditions on the moment conditions, and bounds on the eigenvalues of the weighting matrix so that bounded weights are typically a side effect of these conditions.

A prior with compact support and a move-one-at-a-time proposal are enough to ensure that the Metropolis-Hastings part of the particle Gibbs algorithm will mix (Gamerman and Lopes, 2006). □

4.1 Comments on Particle Filter Performance

The performance of the particle filter depends upon the variance of the weights. For small t there are few degrees of freedom for computing the weighting matrix and the variance of

the weights is a problem. One might try to control this by setting T_0 larger than strictly necessary at the initialization step of the particle filter but doing this has a deleterious effect on the performance of the algorithm because the information from $y_{1:T}$ is not being used until t exceeds T_0 .

A better approach is regularization of the weighting matrix. If the condition number of the weighting matrix (ratio of smallest singular value to the largest) falls below a preset value η (e.g. $\eta = 10^{-8}$) an amount δ is added to the diagonal elements of the weighting matrix just sufficient to bring the condition number to η prior to inversion of the weighting matrix.

Regarding the number of particles one should use in the conditional particle filter, we found that $N = 1000$ gave about the same results as $N = 5000$ and larger. Andrieu, Doucet, and Holenstein (2010) report similar experience for their examples and suggest that the length of the MCMC chain R be increased rather than N because runtimes increase less with R than with N for most of their examples. Because our runtimes increase at the rate $RM[(T!)N + TK]$, the suggestion that N be kept small at the cost of increasing R carries considerable force.

5 Examples

We illustrate our method with two examples: a stochastic volatility model and a DSGE model. In both cases the measurement density is known and thus the examples will provide some insight into information loss and the effect of moment selection in comparison to full information methods. We set $\Psi = \Phi$ so that (9) and (10) become

$$p^*(y, x, \theta) = p^*(y, x | \theta)p^o(x, \theta)p^o(\theta) \quad (15)$$

$$p^*(y | x, \theta) = (2\pi)^{-M/2} \exp\left\{-\frac{1}{2}g_T(y, x, \theta)' [\Sigma(y, x, \theta)]^{-1} g_T(y, x, \theta)\right\}. \quad (16)$$

We use flat priors for $p^o(\theta)$ in order to enable comparison with maximum likelihood estimation and to try to insulate comparisons from the influence of a prior.

5.1 A Stochastic Volatility Model

Our first example is a stochastic volatility (SV) model:

$$y_t = \rho y_{t-1} + \exp(x_t) u_t$$

$$x_t = \phi x_{t-1} + \sigma e_t$$

$$e_t \sim N(0, 1)$$

$$u_t \sim N(0, 1)$$

The true values of the parameters are

$$\theta_0 = (\rho_0, \phi_0, \sigma_0) = (0.9, 0.9, 0.5)$$

for the purpose of plotting the particle filter (PF) and

$$\theta_0 = (\rho_0, \phi_0, \sigma_0) = (0.25, 0.8, 0.1)$$

for illustrating estimation results. The reason for the difference is that the former generates plots that are easy to assess visually whereas the latter is more representative of, say, daily S&P 500 closing prices.

We use the following moment conditions:

$$g_1 = (y_t - \rho y_{t-1})^2 - [\exp(x_t)]^2 \quad (17)$$

$$g_2 = |y_t - \rho y_{t-1}| |y_{t-1} - \rho y_{t-2}| - \left(\frac{2}{\pi}\right)^2 \exp(x_t) \exp(x_{t-1}) \quad (18)$$

\vdots

$$g_{L+1} = |y_t - \rho y_{t-1}| |y_{t-L} - \rho y_{t-L-1}| - \left(\frac{2}{\pi}\right)^2 \exp(x_t) \exp(x_{t-L}) \quad (19)$$

$$g_{L+2} = y_{t-1}(y_t - \rho y_{t-1}) \quad (20)$$

$$g_{L+3} = x_{t-1}(x_t - \phi x_{t-1}) \quad (21)$$

$$g_{L+4} = (x_t - \phi x_{t-1})^2 - \sigma^2 \quad (22)$$

Moment (20) identifies ρ independently of x_t ; moments (17) through (20) overidentify x_t given ρ . Moment (21) identifies ϕ given x_t and moment (22) identifies σ given x_t and ϕ .

What may not be obvious here is how an equation such as (17) identifies x_t . One can see this at the point at which one computes weights in the importance sampling step of the

particle filter algorithm (Section 3). The weight w_t depends on x_t while the weight w_{t-1} does not. Therefore the incremental information regarding x_t provided by (17) does get used at time t to determine x_t . For the particle Gibbs algorithm itself, the incremental information is used at the particle step but is not used at the Metropolis step because the Metropolis step uses sums over all the data rather than partial sums.

Estimates of θ for the SV model are shown in Table 1 for our particle Gibbs GMM method in comparison with the Flury and Shephard (2011) approach. The Flury and Shephard method can be regarded as state-of-the-art. The MCMC chain generated using the method are draws from the exact posterior with a flat prior.

Applying the particle filter at the true value of θ and $N = 5000$, we obtain the estimate of x shown as Figures 1 and 2. The plots for the Flury and Shephard estimator are Figures 3 and 4. In the particle filter vernacular, the particle Gibbs GMM estimator is computed from a smooth whereas the Flury and Shephard estimator is computed from a filter; accordingly, the plots shown for the particle Gibbs GMM estimator are smooths whereas the plots shown of the Flury-Shephard estimator are filters.

(Table 1 about here)

(Figure 1 about here)

(Figure 2 about here)

(Figure 3 about here)

(Figure 4 about here)

5.2 A Dynamic Stochastic General Equilibrium Model

The second example is taken from Del Negro and Schorfheide (2008). We need to have a model with an exact analytical solution to generate accurate data with which to test our proposed method. The working paper version of the article has some simplified versions of the full model that have an analytic expression for the solution. The example is one of the simplified versions.

The full model is a medium-scale New Keynesian model with price and wage rigidities, capital accumulation, investment adjustment costs, variable capital utilization, and habit formation. The simplified model discussed here is obtained by removing capital, fixed costs, habit formation, the central bank, and making wages and prices flexible. With these choices, the model has three shocks: the log difference of total factor productivity z_t , a preference shock that affects intertemporal substitution between consumption and leisure ϕ_t , and the price elasticity of intermediate goods λ_t , called a mark-up shock in the article. In the full model the endogenous variables are output, consumption, investment, capital, and the real wage, which are detrended by $\exp(z_t)$ and expressed as log deviations from the steady-state solution of the model, and inflation. Of these, the ones of interest in the simplified model are the log deviations of wages and output, w_t and y_t , respectively, and inflation π_t . The time increment is one quarter.

The exogenous shocks are

$$\begin{aligned} z_t &= \rho_z z_{t-1} + \sigma_z \epsilon_{z,t} \\ \phi_t &= \rho_\phi \phi_{t-1} + \sigma_\phi \epsilon_{\phi,t} \\ \lambda_t &= \rho_\lambda \lambda_{t-1} + \sigma_\lambda \epsilon_{\lambda,t}, \end{aligned} \tag{23}$$

where $\epsilon_{z,t}$, $\epsilon_{\phi,t}$, and $\epsilon_{\lambda,t}$ are independent standard normal random variables.

The first order conditions are

$$\begin{aligned} 0 &= y_t + \frac{1}{\beta} \pi_t - \mathcal{E}_t(y_{t+1} + \pi_{t+1} + z_{t+1}) \\ 0 &= w_t + \lambda_t \\ 0 &= w_t - (1 + \nu)y_t - \phi_t \end{aligned} \tag{24}$$

where ν is the inverse Frisch labor supply elasticity and β is the subjective discount rate.

The solution for the endogenous variables is

$$\begin{aligned} w_t &= -\lambda_t \\ y_t &= -\frac{1}{1 + \nu} \lambda_t - \frac{1}{1 + \nu} \phi_t \\ \pi_t &= \beta \frac{1 - \rho_\lambda}{(1 + \nu)(1 - \beta \rho_\lambda)} \lambda_t + \beta \frac{1 - \rho_\phi}{(1 + \nu)(1 - \beta \rho_\phi)} \phi_t + \beta \frac{\rho_z}{(1 - \beta \rho_z)} z_t \end{aligned} \tag{25}$$

The true values of the parameters are

$$\theta = (\rho_z, \rho_\phi, \rho_\lambda, \sigma_z, \sigma_\phi, \sigma_\lambda, \nu, \beta) = (0.15, 0.68, 0.56, 0.71, 2.93, 0.11, 0.96, 0.996)$$

which are the parameter estimates for model \mathcal{P}_S of Del Negro and Schorfheide (2008) as supplied by Frank Schorfheide in an email communication.

We take w_t , y_t , and π_t as measured and z_t and ϕ_t as latent.

This model is simple enough that an analytical expression for the likelihood is immediately available by substituting equations (23) into equations (25). By inspection one can anticipate identifications issues: a small change in σ_ϕ can be compensated by small changes to ν , β , and σ_z . This, in turn, causes the MCMC chain for estimating the model by maximum likelihood (Chernozhukov and Hong, 2003) to fail to mix. If one is going to estimate this model by maximum likelihood, one must, as a practical matter, calibrate three of the four parameters σ_z , σ_ϕ , ν , and β . Our choice is to calibrate σ_z , σ_ϕ , and ν , leaving β as the free parameter. The situation here is rather stark: without calibrating σ_z , σ_ϕ , and ν , the MCMC chain for maximum likelihood estimation (MLE) will not mix. Given that the MLE MCMC chain will not mix without these calibrations, one would hardly expect the particle Gibbs GMM chain to mix without them. Indeed, our experience confirms this conjecture.

We shall follow the general principles guiding moment selection set forth in Subsection 2.1, which are to identify as many parameters as possible from the observed data and try to identify the latent variables themselves indirectly from quantities that can be identified from the observed data. The moment conditions (26) – (34) that follow were designed with these principles in mind.

$$g_1 = (w_t - \rho_\lambda w_{t-1})^2 - \sigma_\lambda^2 \tag{26}$$

$$g_2 = w_{t-1}(w_t - \rho_\lambda w_{t-1}) \tag{27}$$

$$g_3 = [w_{t-1} - (1 + \nu)y_{t-1}][w_t - (1 + \nu)y_t - \rho_\phi(w_{t-1} - (1 + \nu)y_{t-1})] \tag{28}$$

$$g_4 = [w_{t-1} - (1 + \nu)y_{t-1}](\phi_t - \rho_\phi \phi_{t-1}) \tag{29}$$

$$g_5 = [w_t - (1 + \nu)y_t]^2 - \sigma_\phi^2 \tag{30}$$

$$g_6 = w_{t-1}(y_{t-1} + \frac{1}{\beta}\pi_{t-1} - y_t - \pi_t - \rho_z z_{t-1}) \tag{31}$$

$$g_7 = y_{t-1}(y_{t-1} + \frac{1}{\beta}\pi_{t-1} - y_t - \pi_t - \rho_z z_{t-1}) \tag{32}$$

$$g_8 = \pi_{t-1}(y_{t-1} + \frac{1}{\beta}\pi_{t-1} - y_t - \pi_t - \rho_z z_{t-1}) \quad (33)$$

$$g_9 = (y_{t-1} + \frac{1}{\beta}\pi_{t-1} - y_t - \pi_t)^2 - \frac{\rho_z^2 \sigma_z^2}{1 - \rho_z^2} \quad (34)$$

Conditions (26) and (27) identify ρ_λ and σ_λ . Recalling that ν is calibrated, (28) identifies ρ_ϕ ; (29) identifies ϕ_t given ρ_ϕ . (This is not literally true because ϕ_t and ρ_ϕ will interact in the Metropolis iterations; this qualification applies a few times below also.) Because both ν and σ_ϕ are calibrated, (30) helps enforce an identity linking w_t and y_t . Because σ_z is calibrated, (31) – (33) identify ρ_z , β , and z_t ; here we cannot identify ρ_z and β without making use of the latent variable z_t , which is likely to negatively affect GMM relative to MLE. However, (34) does help identify ρ_z and β without using z_t .

One could attempt a comparison with the methods proposed in (Fernandez-Villaverde and Rubio-Ramirez, 2006) using equations (25) to avoid numerical solution methods. The difficulty is that (25) is a singular set of measurement equations. The customary approach is to add measurement error to these equations. This presents the additional difficulty of determining how to calibrate the scale of the measurement error. The scale can be manipulated to make results nearly the same as for the MLE (larger scale) or very poor (smaller scale). We do not present these results because we feel one learns nothing from them. One of the advantages of GMM, SMM, and EMM type methods is that singular measurement equations do not cause problems.

Applying the proposed particle Gibbs GMM method to the DSGE model of Subsection 5.2, we obtain the estimates of θ shown in Table 2. Table 2 suggests that the particle Gibbs GMM estimates are reasonable relative to MLE estimates and within the range one might expect for GMM estimates.

(Table 2 about here)

As mentioned in Subsection 2.1, while the moment conditions (26) through (34) can be expected to obtain reasonable results for estimating the parameters θ , they can be expected to do a poor job of estimating the latent variables x . That this is the case here can be verified by inspecting figures similar to Figures 5 and 6 that are not shown. In particular, the plots not shown have slopes that are much shallower than those of Figure 5 and 6.

In order to improve the estimate of x given y we consider the following additional moment conditions derived from the first order conditions of the DSGE model:

$$h_1 = y_{t-1} + \frac{1}{\beta}\pi_{t-1} - y_t - \pi_t - \rho_z z_{t-1} \quad (35)$$

$$h_2 = w_{t-1} h_1 \quad (36)$$

$$h_3 = y_{t-1} h_1 \quad (37)$$

$$h_4 = \pi_{t-1} h_1 \quad (38)$$

$$h_5 = w_t - (1 + \nu)y_t - \phi_t \quad (39)$$

$$h_6 = w_{t-1} h_5 \quad (40)$$

$$h_7 = y_{t-1} h_5 \quad (41)$$

$$h_8 = \pi_{t-1} h_5 \quad (42)$$

Applying the particle filter using conditions (35) through (42) at the true value of θ and $N = 10000$, we obtain the estimates of x given y shown as time series plots in Figure 5 and as scatter plots in Figure 6.

(Figure 5 about here)

(Figure 6 about here)

Estimation results using moment conditions (26) through (34) at the Metropolis step and conditions (35) through (42) at the particle step are shown in Table 3. As seen, by comparing Table 2 to Table 3, estimation performance only improves marginally.

(Table 3 about here)

Using moment conditions (26) through (34) at the Metropolis step and conditions (35) through (42) at the particle step rather than conditions (26) through (34) for both reduces computational cost slightly because runtimes for the particle step increase at approximately $RM(T!)N$ whereas runtimes for the Metropolis step increase at approximately $RMTK$.

A referee’s remark elegantly summarizes the notion of splitting moments: “all moment conditions may be considered present but only (35) through (42) show up when performing the smoothing step.”

5.3 Discussion of Examples

There are several conclusions we can draw from these examples. As expected, in a state space model where an analytic form for the measurement density is available, conventional Bayes when possible, or Flury and Shephard (2011) when not, are better than what we propose unless one is incredibly clever at choosing moment equations. On the other hand, when there is no alternative that does not rely on perturbation or numerical approximations that one would rather avoid, our proposal is a viable option.

The “quality” of the moments matters and there are some principles guiding selection. For the Metropolis algorithm to estimate the parameters θ accurately one should identify as many parameters as possible from the observed data making the latent variables depend as much as possible on quantities that can be computed from the observed data. To track the trajectory of the unobserved latent variables $x_{1:T}$ accurately—the particle part of Algorithm 3—one should choose moments for the particle filter so that observed variables depend on the latent variables as much as possible without regard for identification of parameters. For computational tractability we used two set of moments in the different part of Algorithm 3. A comparison of the results in Table 2 and in Table 3 suggests that computational gains do not come at the expense of precision in the estimation of the static parameter.

However, as noted earlier, economic considerations can render discussion of moment selection moot by dictating the moments that must be used as in the following application.

6 Application

We apply our method to estimate the latent endowment process compatible with a Lucas’ (1978) economy with constant relative risk aversion (CRRA) utility, where we only assume knowledge of the first order conditions of the agents’ optimization problem and an ARCH specification for the latent process. There is currently no other Bayesian method to estimate the latent endowment process without imposing additional assumptions.

The endowment process in a Lucas (1978) economy is typically assumed to equal consumption, whereas our estimates of the latent log endowment growth process indicate that it has larger mean and variance than log consumption growth and we find some evidence of

stochastic volatility.

The agent's first order conditions in a Lucas (1978) economy are

$$1 = \mathcal{E}(M_{t+1}R_{t+1} | \mathcal{F}_t), \quad (43)$$

where $M_{t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^\gamma$ is the marginal rate of substitution, t is time in annual increments, β is the discount factor, γ is the risk aversion parameter, C_t is the endowment process, $R_{t+1} = (P_{t+1} + D_{t+1})/P_t$ is the gross return at time $t + 1$ to an asset that costs P_t at time t with payoff $P_{t+1} + D_{t+1}$ inclusive of dividends and/or interest and \mathcal{F}_t is the agent's information set at time t . We set $\beta = 0.98$ and $\gamma = 2$ as in Mehra and Prescott (1985).

In empirical work C_t is typically measured by annual consumption of nondurables and services and its distribution is usually specified in terms of log consumption growth $x_t = \log(C_t/C_{t-1})$ rather than C_t itself. Here we assume that the endowment process x_t is latent and that it follows an ARCH model:

$$x_t = \mu + \rho x_{t-1} + \sqrt{v_{t-1}} z_t \quad (44)$$

$$v_{t-1} = \sigma^2 + [\tau(x_{t-1} - \mu - \rho x_{t-2})]^2 \quad (45)$$

where z_t is a standard normal, the scores for which are

$$g_1 = \frac{\partial \log f}{\partial z_t} \frac{\partial z_t}{\partial \mu} + \frac{\partial \log f}{\partial \sqrt{v_{t-1}}} \frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} \frac{\partial v_{t-1}}{\partial \mu} \quad (46)$$

$$g_2 = \frac{\partial \log f}{\partial z_t} \frac{\partial z_t}{\partial \rho} + \frac{\partial \log f}{\partial \sqrt{v_{t-1}}} \frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} \frac{\partial v_{t-1}}{\partial \rho} \quad (47)$$

$$g_3 = \frac{\partial \log f}{\partial z_t} \frac{\partial z_t}{\partial \sqrt{v_{t-1}}} \frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} \frac{\partial v_{t-1}}{\partial \sigma} + \frac{\partial \log f}{\partial \sqrt{v_{t-1}}} \frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} \frac{\partial v_{t-1}}{\partial \sigma} \quad (48)$$

$$g_4 = \frac{\partial \log f}{\partial z_t} \frac{\partial z_t}{\partial \sqrt{v_{t-1}}} \frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} \frac{\partial v_{t-1}}{\partial \tau} + \frac{\partial \log f}{\partial \sqrt{v_{t-1}}} \frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} \frac{\partial v_{t-1}}{\partial \tau} \quad (49)$$

where

$$\begin{aligned} \frac{\partial z_t}{\partial \mu} &= -1/\sqrt{v_{t-1}} \\ \frac{\partial z_t}{\partial \rho} &= -x_{t-1}/\sqrt{v_{t-1}} \\ \frac{\partial z_t}{\partial \sqrt{v_{t-1}}} &= -(x_t - \mu - \rho x_{t-1})/v_{t-1} \end{aligned}$$

$$\begin{aligned}
\frac{\partial \sqrt{v_{t-1}}}{\partial v_{t-1}} &= 0.5/\sqrt{v_{t-1}} \\
\frac{\partial v_{t-1}}{\partial \sigma} &= 2\sigma \\
\frac{\partial v_{t-1}}{\partial \tau} &= 2\tau(x_{t-1} - \mu - \rho x_{t-2})^2 \\
\frac{\partial v_{t-1}}{\partial \mu} &= -2\tau^2((x_{t-1} - \mu - \rho x_{t-2})) \\
\frac{\partial v_{t-1}}{\partial \rho} &= -2\tau^2(x_{t-1} - \mu - \rho x_{t-2})x_{t-2} \\
\frac{\partial \log f}{\partial z_t} &= -(x_t - \mu - \rho x_{t-1})/\sqrt{v_{t-1}} \\
\frac{\partial \log f}{\partial \sqrt{v_{t-1}}} &= -1/\sqrt{v_{t-1}}
\end{aligned}$$

We shall reverse engineer the endowment process using Bayesian statistical methods that presume only equations (43) and (44) and standard time series regularity conditions.

Note in particular that M_t is endogenous so that one needs the results of Gallant (2016a) to infer a measurement density from moment conditions that can be used in Bayesian inference. Specifically, to condition on M_t in order to obtain a measurement density one must follow Gallant's protocol to stay within the Bayesian paradigm. Gallant (2016b) contains a detailed discussion of this point.

We compare our results to Gallant and Hong (2007), who treat the marginal rate of substitution M_1, \dots, M_T as an unknown parameter and estimate it nonparametrically using a prior derived from a Bansal and Yaron (2004) economy. We shall use their annual panel data comprised of returns on twenty-four Fama-French (1993) portfolios and thirty-day U.S. Treasury obligations over the period 1930–2004 (one of the twenty-five Fama-French portfolios is lost to missing values). These returns are augmented using information available to agents when portfolios are formed. Specifically, the returns are interacted with a constant, lagged returns on the Fama-French portfolios, lagged debt returns, lagged consumption growth, and lagged labor income growth. All data are real and per capita. For a discussion of the ideas involved in this augmentation see Gallant, Hansen, and Tauchen (1990) and Hansen and Jagannathan (1991). For a discussion of the factors relevant to pricing payoffs and the role that labor income growth plays see Campbell (1996). For the specific details of data

construction see Gallant and Hong (2007). The observable variables in our application are returns, consumption growth, and labor income growth and we denote them by y .

The differences between the nonparametric approach of Gallant and Hong (2007) and this paper can be seen from their respective posteriors

$$p_{gh}(y | \theta_{(1)}) p_{gh}(\theta_{(1)} | \theta_{(2)}) p_{gh}(\theta_{(2)}) \quad (50)$$

$$p_{gh}\{y | \beta[\exp(x)]^\gamma\} p_{ggr}(x | \theta_{(3)}) p_{ggr}(\theta_{(3)}), \quad (51)$$

where $\theta_{(1)}, \theta_{(2)}, \theta_{(3)}$ refer to different sets of parameters. The functional form of the density $p_{gh}(\cdot | \cdot)$ is the same in both (50) and (51); it is the normal density $\phi(z)$ with Z given by (7) using (54) below as the moments g in (7).⁵

In Gallant and Hong there are no latent variables: Inference is by means of straightforward MCMC using a move-one-at-a-time proposal density on the parameters $(\theta_{(1)}, \theta_{(2)})$.

Let s_t denote the vector of gross returns at time t on the twenty-four Fama-French portfolios and let b_t denote the gross returns at time t on the thirty-day Treasury debt issue. Let $c_t = C_t/C_{t-1} = \exp(x_t)$ denote endowment growth and let l_t denote labor income growth. Define the instruments

$$V_t = \begin{pmatrix} s_t - 1 \\ b_t - 1 \\ c_t - 1 \\ l_t - 1 \\ 1 \end{pmatrix},$$

where $s_t - 1$ and $b_t - 1$ denote 1 subtracted from each element of s_t and b_t . Denote the vector of Euler equation errors by

$$e_t(s_{t+1}, b_{t+1}, M_{t+1}) = 1 - M_{t+1} \begin{pmatrix} s_{t+1} \\ b_{t+1} \end{pmatrix}, \quad (52)$$

where 1 denotes a vector of 1's of length twenty-seven. Consider the moment functions

$$m_t(s_t, b_t, c_t, l_t, s_{t+1}, b_{t+1}, M_{t+1}) = V_t \otimes e_t(s_{t+1}, b_{t+1}, M_{t+1}), \quad (53)$$

⁵There are defects in the derivation of $p_{gh}(\cdot | \cdot)$ in Gallant and Hong that are corrected in Gallant (2016a, 2016b). Therefore, in this paper, we are relying on Section 2 for the correct derivation of $p_{gh}(\cdot | \cdot)$.

where $t = 1, \dots, T = 75$. The length of the vector m_t is 700. One can view (53) as a set of payoffs that has been enlarged by interacting returns with V_t so that the actual set of payoffs under consideration is $V_t \otimes (s_{t+1}, b_{t+1})$. On this see Gallant, Hansen, and Tauchen (1990) and Hansen and Jagannathan (1991). Following Gallant and Hong (2007) we assume that $(M_t s_t, M_t b_t)$ has a factor structure: There is one error common to all elements of $M_t s_t$ and twenty-four idiosyncratic errors, one for each element of $M_t s_t$. Denote this matrix by Σ_e . A set of orthogonal eigenvectors U_e for Σ_e are easy to construct and can be used to diagonalize Σ_e . To illustrate, if there were four stocks, then

$$U_e = \begin{pmatrix} 1/4 & 1/\sqrt{2} & 1/\sqrt{6} & 1/\sqrt{12} & 0 \\ 1/4 & -1/\sqrt{2} & 1/\sqrt{6} & 1/\sqrt{12} & 0 \\ 1/4 & 0 & -2/\sqrt{6} & 1/\sqrt{12} & 0 \\ 1/4 & 0 & 0 & -3/\sqrt{12} & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Similarly U_v and Σ_v for V_t . (U_v looks like U_e but with an extra 2x2 block for (c_t, l_t) and a one appended to the southeast corner.) Then equation (4) becomes

$$g_T(y, x, \theta) = \frac{1}{\sqrt{T}} (U_v \otimes U_e)' m_t(s_t, b_t, c_t, l_t, s_{t+1}, b_{t+1}, M_{t+1}), \quad (54)$$

where $y_{t+1} = (s_t, b_t, c_t, l_t, s_{t+1}, b_{t+1})$, x_{t+1} is given by (44), and $\theta = (\mu, \rho, \sigma, \tau, \beta, \gamma)$. Expression (5) is used to compute the weighting matrix.

Our prior for θ is a truncated normal determined as follows. Annual, real, per capita, U.S. log consumption growth data approximately follows (44) with $\tau = 0$ and $(\mu, \rho, \sigma) = (0.015, 0.35, 0.015)$ over most subperiods of the dates 1925 through 2015. We choose these values of (μ, ρ, σ) for the location parameter of the normal part of our prior and the location parameter for τ to 0.01 to shift τ somewhat to the right of zero. Scale parameters are set so that the marginal probabilities of being within $S \times 100\%$ of the location parameter is 95%. The joint is the product of the marginals. The support conditions are $-100 < \mu < 100$, $-.999 < \rho < .999$, $0 < \sigma < 100$, and $0 < \tau < 100$, which more than encompasses all reasonable values of these four parameters. We computed estimates using scale factors $S=50, 500$, and 1000 and report them in Table 4. The prior for (β, γ) is $(\beta, \gamma) = (0.98, 2)$ with probability one. The proposal for the MCMC chain is move-one-at-a-time random walk.

Estimation results are presented in Table 4. They suggest that location and scale of estimated log endowment growth shift upwards relative to log consumption growth and there is evidence of stochastic volatility. The evidence is not overwhelming because the standard deviations of the posterior are large. Converted to the marginal rate of substitution M_t , the estimated endowment process and the Gallant and Hong (2009) estimate of the MRS are plotted in Figure 7 as a time series and in Figure 8 as a scatter plot. Relative to the results of Gallant and Hong, the estimated marginal rate of substitution is more volatile because, apparently, the prior used by Gallant and Hong is tighter than our prior here. Both suggest the presence of jumps, but the evidence here is stronger. Histograms of draws from the posterior for the parameters of the log endowment process are displayed in Figure 9. The distributions are rather markedly skewed.

(Table 4 about here)

(Figure 7 about here)

(Figure 8 about here)

(Figure 9 about here)

7 Conclusion

We proposed an algorithm for Bayesian estimation of the parameters of a dynamic model with latent dynamic variables when the model does not provide a measurement density but only a set of moment conditions involving observable and latent variables. The algorithm is a modification of a particle filter algorithm where the measurement density is substituted with its “GMM representation”. We showed how to construct such a density and provided a theoretical justification. We illustrated with two examples: a stochastic volatility model and a dynamic stochastic general equilibrium model. We applied the method to estimating the latent endowment process compatible with a Lucas (1978) economy and found that it has dramatically different properties than the consumption time series that are typically used in the literature.

8 References

- Andrews, Donald W. K., 1991, Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica* 59, 817–858.
- Andrieu, C., A. Doucet, and R. Holenstein, 2010, Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B* 72, 269–342.
- Andrieu C. and Roberts, G. O., 2009, The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics* 37, 697–725.
- Bansal, R., and A. Yaron. 2004. Risks for the long run: a potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.
- Benigno, G., Benigno, P., and Nisticó, S., 2012, Risk, monetary policy and the exchange rate. *NBER Macroeconomics Annual* 26, 247–309.
- Beaumont, M. A., 2003, Estimation of population growth or decline in genetically monitored populations. *Genetics* 164 1139–1160.
- Bloom, Nicholas, 2009, The impact of uncertainty shocks. *Econometrica* 77, 623–685.
- Campbell, J. Y., 1996, Understanding risk and return. *The Journal of Political Economy* 104, 298–345.
- Chernozhukov, Victor, and Han Hong, 2003, An MCMC approach to classical estimation. *Journal of Econometrics* 115, 293–346.
- Chopin, Nicolas; Singh, Sumeetpal S, 2015, On particle Gibbs sampling. *Bernoulli* 21, 1855–1883.
- Cochrane, John H., 2005, *Asset pricing (Revised Edition)*. Princeton University Press, Ch. 16, esp. pp. 292 & 298.
- Del Negro, Marco, and Frank Schorfheide, 2008, Forming priors for DSGE models (and how it affects the assessment of nominal rigidities). *Journal of Monetary Economics* 55, 1191–1208.

Fama, E., and K. French., 1993, Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3–56.

Fernandez-Villaverde, J., and J. F. Rubio-Ramirez, 2006, Estimating macroeconomics models: a likelihood approach. NBER Technical Working Paper No. 321. <http://www.nber.org/papersby>

Flury, Thomas, and Neil Shephard, 2011, Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models. *Econometric Theory* 27, 933–956.

Gallant, A. Ronald, 2016a, Reflections on the probability space induced by moment conditions with implications for Bayesian inference. *Journal of Financial Econometrics* 14, 227–228.

Gallant, A. Ronald, 2016b, Reply to comments on Reflections. *Journal of Financial Econometrics* 14, 284–294.

Gallant, A. Ronald, 2016c, Addendum to Reply to Reflections. Working paper, <http://www.aronaldg.org>

Gallant, A. Ronald, 2016d, Complementary Bayesian Method of Moments Strategies. Working paper, <http://www.aronaldg.org>.

Gallant, A. Ronald, and Han Hong, 2007, A statistical inquiry into the plausibility of recursive utility. *Journal of Financial Econometrics* 5, 523–590.

Gallant, A. Ronald, Han Hong, and Ahmed Khwaja, 2015, A Bayesian approach to estimation of dynamic models with small and large number of heterogeneous players and latent serially correlated states. Working paper, <http://www.aronaldg.org>.

Gallant, A. R., L. P. Hansen, and G. Tauchen, 1990, Using conditional moments of asset payoffs to infer the volatility of intertemporal marginal rates of substitution. *Journal of Econometrics* 45, 141–180.

Gallant, A. R., and R. E. McCulloch, 2009, On the determination of general statistical models with application to asset pricing. *Journal of the American Statistical Association* 104, 117–131.

- Gallant, A. Ronald, and Douglas W. Nychka, 1987, Semi-nonparametric maximum likelihood estimation. *Econometrica* 55, 363–390.
- Gallant, A. Ronald, and George Tauchen, 1996, Which moments to match? *Econometric Theory* 12, 657–681.
- Gallant, A. Ronald, and George Tauchen, 2015, Efficient indirect Bayesian estimation: with application to the dominant generalized Blumenthal-Gettoor index of an Ito semimartingale. Working paper, www.aronaldg.org.
- Gallant, A. Ronald, and Halbert L. White, Jr., 1987, A unified theory of estimation and inference for nonlinear dynamic models. Basil Blackwell Ltd., Oxford. <http://www.aronaldg.org/webfile>
- Gamerman, D., and H. F. Lopes, 2006, Markov chain Monte Carlo: stochastic simulation for Bayesian inference (2nd edition). Chapman and Hall, Boca Raton, FL.
- Hansen, L. P., 1982, Large sample properties of generalized method of moments estimators. *Econometrica* 50, 1029–1054.
- Hansen, L. P., and R. Jagannathan, 1991, Implications of security market data for models of dynamic economics. *Journal of Political Economy* 99 225–262.
- Hansen, L. P., and Singleton, K. J., 1982, Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50 1269–1286.
- Lucas, R. E., Jr., 1978, Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.
- Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: a puzzle. *Journal of Monetary Economics* 15, 145–161.
- Pitt, M. K., and N. Shephard, 1999, Filtering via simulation: auxiliary particle filters. *Journal of the American statistical association* 94, 590–599.
- Sweeting, T.J., 1986, On a converse to Scheffe’s theorem. *The Annals of Statistics* 14, 1252–1256.

Algorithm 1 Particle filter with GMM representation of the measurement density

Step 1. Initialization.

- (a) Given θ (and $y_{1:T}$),
- (b) Set T_0 to the smallest t required to compute $g_t(y_{1:t}, x_{1:t}, \theta)$ and $\Sigma(y_{1:t}, x_{1:t}, \theta)$.
- (c) For $i = 1, \dots, N$ sample $(x_1^i, x_2^i, \dots, x_{T_0}^i)$ from $p^o(x_t|x_{t-1}, \theta)$.
- (d) Set t to $T_0 + 1$.
- (e) Set $x_{1:t-1}^i = (x_1^i, x_2^i, \dots, x_{T_0}^i)$.

Step 2. Importance sampling.

- (a) For $i = 1, \dots, N$ sample \tilde{x}_t^i from $p^o(x_t|x_{t-1}, \theta)$ and set

$$\tilde{x}_{1:t}^i = (x_{0:t-1}^i, \tilde{x}_t^i).$$

- (b) For $i = 1, \dots, N$ compute weights $w_t^i(\theta) = p^*(y_{1:t} | \tilde{x}_{1:t}^i, \theta)$.
- (c) Scale the weights to sum to one,

$$W^i = \frac{w^i(\theta)}{\sum_{i=1}^N w^i(\theta)}.$$

Step 3. Selection.

- (a) For $i = 1, \dots, N$ sample with replacement particles $x_{1:t}^i$ from the set $\{\tilde{x}_{1:t}^i\}$ according to the weights $\{W^i\}_{i=1}^N$.

Step 4. Repeat

- (a) If $t < T$, increment t and go to importance sampling step;
 - (b) else output $\{x_{1:T}^i\}_{i=1}^N$.
-

Algorithm 2 PMMH with GMM representation of the measurement density

Step 1 Initialization, $i = 1$.

- (a) Set $\theta^{(0)}$ arbitrarily.
- (b) Run the particle filter of Algorithm 1, obtain a sample $x_{1:T}^{(0)}$ from $p^*(\cdot|y_{1:T}|\theta^{(0)})$ and let $\hat{p}^*(y_{1:T}, \theta^{(0)})$ be the marginal distribution estimate (14).

Step 2 for $i > 1$,

- (a) Sample $\theta^{prop} \sim q(\cdot|\theta^{(i-1)})$
- (b) Run the particle filter of Algorithm 1, obtain a sample $x_{1:T}^{prop}$ from $p^*(\cdot|y_{1:T}, \theta^{prop})$ and let $\hat{p}^*(y_{1:T}|\theta^{prop})$ the marginal distribution estimate.
- (c) With probability

$$\min \left\{ 1, \frac{\hat{p}^*(y_{1:T}|\theta^{prop})p(\theta^{prop})}{\hat{p}^*(y_{1:T}|\theta^{(i-1)})p(\theta^{(i-1)})} \frac{q(\theta^{(i-1)}|\theta^{prop})}{q(\theta^{prop}|\theta^{(i-1)})} \right\}$$

set $\theta^{(i)} = \theta^{prop}$, $x_{1:T}^{(i)} = x_{1:T}^{prop}$, and $\hat{p}^*(y_{1:T}|\theta^{(i)}) = \hat{p}^*(y_{1:T}|\theta^{prop})$;

- (d) otherwise set $\theta^{(i)} = \theta^{(i-1)}$, $x_{1:T}^{(i)} = x_{1:T}^{(i-1)}$, and $\hat{p}^*(y_{1:T}|\theta^{(i)}) = \hat{p}^*(y_{1:T}|\theta^{(i-1)})$.
-

Algorithm 3 Particle Gibbs algorithm targeting $p^*(\theta, x_{1:T}|y_{1:T})$.

Step 1 Start

- (a) Set $j = 0$.
- (b) Set $\hat{x}_{1:T}^{(1)}$ and $\hat{\theta}^{(1)}$ from the output of Algorithm 2, from some other estimation scheme, or arbitrarily subject to support conditions.

Step 2 Conditional particle filter

- (a) Initialization.
 - (i) Increment j and set $\tilde{x}_{1:T}^1 = x_{1:T}^1 = \hat{x}_{1:T}^{(j)}$.
 - (ii) Set T_0 to the smallest t required to compute $g_t(y_{1:t}, x_{1:t}, \theta)$ and $\Sigma(y_{1:t}, x_{1:t}, \theta)$.
 - (iii) For $i = \mathbf{2}, \dots, N$ sample $(x_1^i, x_2^i, \dots, x_{T_0}^i)$ from $p^o(x_t|x_{t-1}, \theta)$.
 - (iv) Set t to $T_0 + 1$.
 - (v) Set $x_{1:t-1}^i = (x_1^i, x_2^i, \dots, x_{T_0}^i)$ for $i = \mathbf{2}, \dots, N$.
- (b) Importance sampling step.
 - (i) For $i = \mathbf{2}, \dots, N$ sample \tilde{x}_t^i from $p^o(x_t|x_{t-1}^i, \theta)$ and set $\tilde{x}_{1:t}^i = (x_{0:t-1}^i, \tilde{x}_t^i)$.
 - (ii) For $i = \mathbf{1}, \dots, N$ compute weights $\tilde{w}_t^i = p^*(y_{1:t} | \tilde{x}_{1:t}^i, \theta)$.
 - (iii) Scale the weights to sum to one, $W^i = \tilde{w}^i / \sum_{i=1}^N \tilde{w}^i$.
- (c) Selection step.
 - (i) For $i = \mathbf{2}, \dots, N$ sample with replacement particles $x_{1:t}^i$ from the set $\{\tilde{x}_{1:t}^i\}_{i=1}^N$ according to the weights $\{W^i\}_{i=1}^N$.
- (d) Repeat conditional particle filter
 - (i) If $t < T$, increment t and go to Importance sampling step;
 - (ii) else put $\hat{x}^{(j+1)} = x_{1:T}^N$ and go to Metropolis step.

Step 3 Metropolis step

- (a) Initialize: Put $i = 1$ and $\theta^i = \hat{\theta}^{(j)}$.
- (b) Propose: Draw θ^{prop} from $q(\theta|\theta^i)$
- (c) Accept-Reject: Put θ^{i+1} to θ^{prop} with probability

$$\alpha = \min \left[1, \frac{p^*(y, \hat{x}^{(j+1)}, \theta^{prop})q(\theta^i|\theta^{prop})}{p^*(y, \hat{x}^{(j+1)}, \theta^i)q(\theta^{prop}|\theta^i)} \right]$$

else put θ^{i+1} to θ^i .

- (d) Repeat Metropolis
 - (i) If $i + 1 < K$ increment i and go to Propose;
 - (ii) else set $\hat{\theta}^{(j+1)} = \theta^K$ and go to Repeat outer loop.

Step 4 Repeat outer loop

- (a) If $j + 1 < R$ go to Conditional particle filter step.
- (b) else terminate; $\{\hat{\theta}^{(j)}, \hat{x}^{(j)}\}_{j=1}^R$ is the MCMC chain.

Table 1. Parameter Estimates for the SV Model Using Moment Conditions (17) through (22) at both the Metropolis and Particle Steps.

Parameter	True Value	Mean	Mode	Standard Error
Bayesian				
ρ	0.25	0.30271	0.30939	0.076758
ϕ	0.8	0.15348	0.85765	0.643400
σ	0.1	0.11400	0.08435	0.070081
Flury and Shephard Estimator				
ρ	0.25	0.30278	0.28555	0.059320
ϕ	0.8	0.17599	0.89189	0.509780
σ	0.1	0.09737	0.07839	0.064661

Data of length $T = 250$ was generated by simulating the model of Subsection 5.1 at the parameter values shown in the column labeled “True Value”. In the first panel the model was estimated by using the particle Gibbs method (Algorithm 3) described in Section 3 with a one-lag HAC weighting matrix using $N = 1000$ particles and $K = 50$ draws for Metropolis. In the second panel the estimator is the Bayesian estimator proposed by Flury and Shepard (2010) with a flat prior. It is a standard maximum likelihood particle filter estimator except that the seed changes every time a new θ is proposed with N increased as necessary to control the rejection rate of the MCMC chain. The columns labeled mean, mode, and standard deviation are the mean, mode, and standard deviations of a particle Gibbs chain of length $R = 9637$ for the first panel and the same from an MCMC chain of length $R = 500000$ with a stride of 5 for the second.

Table 2. Parameter Estimates for the DSGE Model Using Moment Conditions (26) through (34) at Both the Metropolis and Particle Steps.

Parameter	True Value	Mean	Mode	Standard Error
Bayesian				
ρ_z	0.15	0.21887	0.23069	0.09179
ρ_ϕ	0.68	0.59967	0.60750	0.04988
ρ_λ	0.56	0.50884	0.31473	0.28981
σ_λ	0.11	0.10797	0.11613	0.06896
β	0.996	0.98201	0.99634	0.01834
Maximum Likelihood				
ρ_z	0.15	0.15165	0.15087	0.00583
ρ_ϕ	0.68	0.59185	0.59419	0.05044
ρ_λ	0.56	0.56207	0.56549	0.05229
σ_λ	0.11	0.11225	0.11189	0.00508
β	0.996	0.99640	0.99643	0.00186

Data of length $T = 250$ was generated by simulating the model of Subsection 5.2 at the parameter values shown in the column labeled “True Value”. In the first panel the model was estimated by using the particle Gibbs method (Algorithm 3) described in Section 3 with a two-lag HAC weighting matrix using $N = 1000$ particles and $K = 50$ draws for Metropolis. In the second panel the model was estimated by maximum likelihood. The columns labeled mean, mode, and standard deviation are the mean, mode, and standard deviations of a particle Gibbs chain of length $R = 9637$ for the first panel and the same from an MCMC chain of length $R = 500000$ with a stride of 5 for the second.

Table 3. Parameter Estimates for the DSGE Model Using Conditions (26) through (34) at the Metropolis Step and Conditions (35) through (42) at the Particle Step

Parameter	True Value	Mean	Mode	Standard Error
Bayesian				
ρ_z	0.15	0.23508	0.15007	0.08975
ρ_ϕ	0.68	0.69870	0.58945	0.06127
ρ_λ	0.56	0.49904	0.46443	0.28418
σ_λ	0.11	0.11292	0.08924	0.06559
β	0.996	0.97465	0.99604	0.02479
Maximum Likelihood				
ρ_z	0.15	0.15165	0.15087	0.00583
ρ_ϕ	0.68	0.59185	0.59419	0.05044
ρ_λ	0.56	0.56207	0.56549	0.05229
σ_λ	0.11	0.11225	0.11189	0.00508
β	0.996	0.99640	0.99643	0.00186

As for Table 2.

Table 4. Parameter Estimates for the MRS Model Using Conditions (46) through (49) at the Metropolis Step and Condition (54) at the Particle Step

Parameter	Mean	Standard Deviation
Prior Scale Factor = 50		
μ	0.18203	0.13172
ρ	0.56149	0.28256
σ	0.48836	0.20040
τ	0.23083	0.28940
β	0.98	0.0
γ	2.0	0.0
Prior Scale Factor = 500		
μ	0.22091	0.19625
ρ	0.51358	0.40154
σ	0.36577	0.34499
τ	1.68040	1.37980
β	0.98	0.0
γ	2.0	0.0
Prior Scale Factor = 1000		
μ	0.18904	0.25181
ρ	0.33079	0.56017
σ	0.29808	0.36470
τ	2.17500	1.75620
β	0.98	0.0
γ	2.0	0.0

The model was estimated from 75 real, per-capita, annual observations from 1930 to 2004 on twenty-four Fama-French (1993) portfolios, thirty-day U.S. Treasury debt obligations, consumption growth, and labor income growth using the particle Gibbs method (Algorithm 3) described in Section 3 with weighting matrix given by (5) using $N = 500$ for the conditional particle step and $K = 50$ draws for the Metropolis step. Parameters β and γ are fixed as shown. The columns labeled mean and standard deviation are the mean and standard deviations of a particle Gibbs chain of length $R = 9610$. For measured log consumption growth, the parameter values are $(\mu, \rho, \sigma, \tau) = (0.015, 0.35, 0.015, 0.0)$ over most subperiods of the dates 1925 through 2015.




Figure 1. PF for Λ , Time Series Plot, SV Model. Data of length $T = 100$ was generated from a simulation of the stochastic volatility model of Subsection 5.1 and $N = 5000$ particles computed using the particle filter algorithm (Algorithm 1) described in Section 3. The dashed blue line plots the simulated Λ . The solid red line is the mean of the particles and the dotted red lines are plus and minus two pointwise standard errors. The moment equations were (17) through (22); a one lag HAC estimator was used for (5).




Figure 2. PF for Λ , Scatter Plot, SV Model. As for Figure 1 except that plotted is the mean of the particles vs. the simulated Λ .




Figure 3. PF for Λ , Flurry-Shephard Method, Time Series Plot, SV Model. As for Figure 1 except that plotted is a filter, not a smooth, and weighting is by the actual measurement density, not the “GMM representation” of the measurement density.




Figure 4. PF for Λ , Flurry-Shephard Method, Scatter Plot. As for Figure 3 except that plotted is the mean of the particles vs. the simulated Λ .




Figure 5. PF for Λ , Time Series Plot, DSGE Model. Data of length $T = 250$ was generated by simulating the DSGE model of Subsection 5.2 and $N = 10000$ particles were computed using the particle filter algorithm (Algorithm 1) described in Section 3. The dashed blue line in the upper panel plots the simulated ϕ_t for the last 50 time points. The lower panel is the same for z_t . In both panels, the solid red line is the mean of the particles and the dotted red lines are plus and minus two pointwise standard errors. The moment equations were (35) through (42); a two lag HAC estimator was used for (5).




Figure 6. PF for Λ , Scatter Plot, DSGE Model. As for Figure 5 except that plotted is the mean of the particles vs. the simulated Λ for all 250 time points.




Figure 7. Particle Filter for MRS, Time Series Plot. The model was estimated as described in Table 4 for the case scale factor 500. The dashed blue line plots the MRS estimated by Gallant and Hong (2007). The solid red line is the mean of the particles transformed to MRS using $M_t = \beta[\exp(x_t)]^\gamma$




Figure 8. Particle Filter for MRS, Scatter Plot. As for Figure 7 except that plotted is the mean of the particles vs. the MRS estimated by Gallant and Hong (2007).

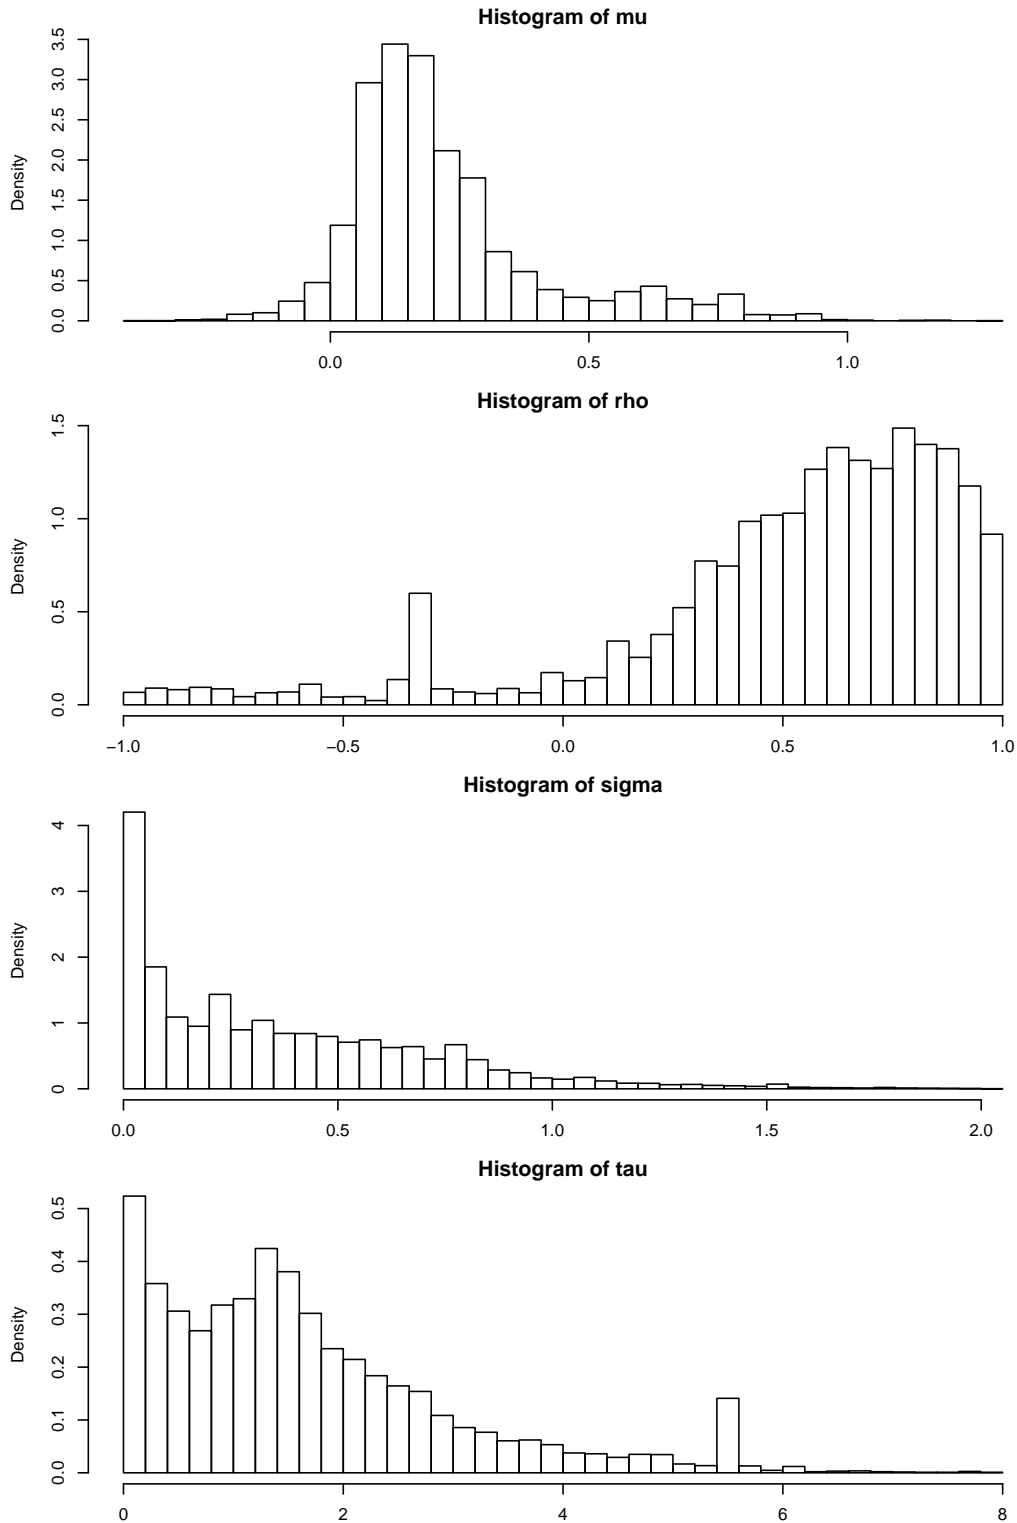


Figure 9. Particle Filter for MRS, Time Series Plot. The figure shows the histograms of the draws in the MCMC chain described in Table 4 for the case scale factor 500. $M_t = \beta[\exp(x_t)]^\gamma$