

STATE-OBSERVATION SAMPLING AND THE ECONOMETRICS OF LEARNING MODELS

Laurent E. Calvet and Veronika Czellar*

First version: February 2010

This version: May 2011

Abstract

In nonlinear state-space models, sequential learning about the hidden state can proceed by particle filtering when the density of the observation conditional on the state is available analytically (e.g. Gordon *et al.* 1993). This condition need not hold in complex environments, such as the incomplete-information equilibrium models considered in financial economics. In this paper, we make two contributions to the learning literature. First, we introduce a new filtering method, the state-observation sampling (SOS) filter, for general state-space models with intractable observation densities. Second, we develop an indirect inference-based estimator for a large class of incomplete-information economies. We demonstrate the good performance of these techniques on an asset pricing model with investor learning applied to over 80 years of daily equity returns.

Keywords: Hidden Markov model, particle filter, state-observation sampling, learning, indirect inference, forecasting, state space model, value at risk.

*Calvet: Department of Finance, HEC Paris, 1 rue de la Libération, 78351 Jouy en Josas, France; calvet@hec.fr. Czellar: Department of Economics and Decision Sciences, HEC Paris, 1 rue de la Libération, 78351 Jouy en Josas, France, czellarv@hec.fr. An earlier version of this paper was circulated under the title “Efficient Estimation of Learning Models.” We received helpful comments from Pavel Chigansky, Nicolas Chopin, Adlai Fisher, Thierry Foucault, Itzhak Gilboa, Christian Gouriéroux, Lars Hansen, Per Mykland, Hashem Pesaran, Nicholas Polson, Elvezio Ronchetti, Andrew Siegel, Ramon van Handel, Paolo Zaffaroni, and seminar participants at CORE, the University of Chicago, the University of Evry, the Second HEC Finance and Statistics Conference, CEF 2010 in London, the 2010 Toulouse School of Economics Financial Econometrics Conference, ISBIS 2010 in Portoroz, and the 2010 Econometric Society World Congress in Shanghai. We gratefully acknowledge the computer support of EUROFIDAI and the financial support of the American Statistical Association, the Europlace Institute of Finance, and the National Institute of Statistical Sciences.

1 Introduction

Sequential learning by economic agents is a powerful mechanism that theoretically explains key properties of asset returns, aggregate performance and other equilibrium outcomes (e.g., Pástor and Veronesi, 2009a).¹ In order to use these models in practice, for instance to forecast and price assets, a crucial question arises: How can we track agent beliefs? A natural possibility is to consider particle filters, a large class of sequential Monte Carlo methods designed to track a hidden Markov state from a stream of partially revealing observations (e.g. Gordon, Salmond, and Smith, 1993; Johannes and Polson, 2009; Pitt and Shephard, 1999). Existing filtering methods, however, are based on the assumption that the density of the observation conditional on the hidden state (called *observation density*) is available in closed form up to a normalizing constant. This assumption is unfortunately not satisfied in incomplete-information economies. In this paper, we introduce the state-observation sampling (SOS) filter, a novel sequential Monte Carlo method for general state space models with intractable observation densities. In addition, we develop an indirect inference-based estimator (Gouriéroux, Monfort and Renault 1993; Smith, 1993) for the structural parameters of an incomplete-information economy.

Since their introduction by Gordon, Salmond, and Smith (1993), particle filters have considerably expanded the range of applications of hidden Markov models and now pervade fields as diverse as engineering, genetics, statistics (Andrieu and Doucet, 2002; Chopin, 2004; Kuensch, 2005), finance (e.g. Kim, Shephard and Chib, 1998; Johannes, Polson, and Stroud, 2009), and macroeconomics (Fernández-Villaverde and Rubio-Ramirez, 2007; Fernandez-Villaverde et al., 2009; Hansen, Polson and Sargent, 2011).² These methods provide estimates of the distribution of a hidden Markov state s_t conditional on a time

¹In financial economics, investor learning has been used to explain phenomena as diverse as the level and volatility of equity prices, return predictability, portfolio choice, mutual fund flows, firm profitability following initial public offerings, and the performance of venture capital investments. In particular, the portfolio and pricing implications of learning are investigated in Brennan (1998), Brennan and Xia (2001), Calvet and Fisher (2007), David (1997), Guidolin and Timmermann (2003), Hansen (2007), Pástor and Veronesi (2009b), Timmermann (1993, 1996), and Veronesi (1999, 2000). We refer the reader to Pástor and Veronesi (2009a) for a recent survey of learning in finance.

²Advances in particle filtering methodology include Andrieu, Doucet, and Holenstein (2010), Del Moral (2004), Fearnhead and Clifford (2003), Gilks and Berzuini (2001), Godsill, Doucet, and West (2004), and Storvik (2002). Particle filters have received numerous applications in finance, such as model diagnostics (Chib, Nardari, and Shephard, 2002), simulated likelihood estimation (Pitt, 2005), volatility forecasting (Calvet, Fisher, and Thompson, 2006), and derivatives pricing (Christoffersen, Jacobs, and Mimouni 2007). See Cappé, Moulines and Rydén (2005), Doucet and Johansen (2008), and Johannes and Polson (2009) for recent reviews.

series of observations $R_t = (r_1, \dots, r_t)$, $r_t \in \mathbb{R}^{n_R}$, by way of a set of “particles” $(s_t^{(1)}, \dots, s_t^{(N)})$. In the original sampling and importance resampling algorithm of Gordon, Salmond, and Smith (1993), the construction of the date- t filter from the date- $(t-1)$ particles proceeds in two steps. In the mutation phase, a new set of particles is obtained by drawing a hidden state $\tilde{s}_t^{(n)}$ from each date- $(t-1)$ particle $s_{t-1}^{(n)}$ under the transition kernel of the Markov state. Given a new observation r_t , the particles are then resampled using weights that are proportional to the observation density $f_R(r_t | \tilde{s}_t^{(n)}, R_{t-1})$. Important refinements of the algorithm include sampling from an auxiliary model in the mutation phase (Pitt and Shephard, 1999), or implementing variance-reduction techniques such as stratified (Kitagawa 1996) and residual (Liu and Chen 1998) resampling.

A common feature of existing filters is the requirement that the observation density $f_R(r_t | s_t, R_{t-1})$ be available analytically up to a normalizing constant. This condition need not hold in economic models in which equilibrium conditions can create complex nonlinear relationships between observations and the underlying state of the economy. In the special case when the state s_t evolves in a Euclidean space \mathbb{R}^{n_S} and has a continuous distribution, a possible solution is to estimate each observation density $f_R(r_t | \tilde{s}_t^{(n)}, R_{t-1})$, $n \in \{1, \dots, N\}$, by nonparametric methods (Rossi and Vila, 2006, 2009). This approach is numerically challenging because N conditional densities, and therefore $2N^2$ kernels, must be evaluated every period. Furthermore, the rate of convergence decreases both with the dimension of the state space, n_S , and the dimension of the observation space, n_R , which indicates that the algorithm is prone to the curse of dimensionality.

The present paper develops a novel particle filter for general state space models that does not require the calculation of the observation density. This new method, which we call the State-Observation Sampling (SOS) filter, consists of simulating a state and a pseudo-observation $(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})$ from each date- $(t-1)$ particle. In the resampling stage, we assign to each particle $\tilde{s}_t^{(n)}$ an importance weight determined by the proximity between the pseudo-observation $\tilde{r}_t^{(n)}$ and the actual observation r_t . We quantify proximity by a kernel of the type considered in nonparametric statistics:

$$p_t^{(n)} \propto \frac{1}{h_t^{n_R}} K \left(\frac{r_t - \tilde{r}_t^{(n)}}{h_t} \right),$$

where h_t is a bandwidth, and K is a probability density function. The resampling stage tends to select states associated with pseudo-observations in the neighborhood of the actual

data. SOS requires the calculation of only N kernels each period and makes no assumptions on the characteristics of the state space, which may or may not be Euclidean. We demonstrate that as the number of particles N goes to infinity, the filter converges to the target distribution under a wide range of conditions on the bandwidth h_t . The root mean squared error of moments computed using the filter decays at the rate $N^{-2/(n_R+4)}$, that is at the same rate as the kernel density estimator of a random vector on \mathbb{R}^{n_R} . The asymptotic rate of convergence is thus invariant to the size of the state space, indicating that SOS overcomes a form of the curse of dimensionality. We also prove that the SOS filter provides consistent estimates of the likelihood function.

We next develop inference methods for incomplete-information equilibrium models. To clarify the exposition, we focus on a class of recursive incomplete-information economies parameterized by $\theta \in \Theta$, which nests the examples of Brandt, Zeng, and Zhang (2004), Calvet and Fisher (2007), David and Veronesi (2006), Lettau, Ludvigson and Wachter (2008), Moore and Schaller (1996) and van Nieuwerburgh and Veldkamp (2006). We consider three levels of information, which correspond to nature, an agent and the econometrician. Figure 1 illustrates the information structure. At the beginning of every period t , nature selects a Markov *state of nature* M_t and a vector of fundamentals or signals x_t , whose distribution is contingent on the state of nature. The agent observes the signal x_t , and computes the conditional probability distribution (“belief”) $\Pi_t = \Pi_t(x_t, \Pi_{t-1})$, for instance by using Bayes’ rule. According to her beliefs and signal, the agent also computes a data point $r_t = \mathcal{R}(x_t, \Pi_t, \Pi_{t-1}; \theta)$, which may for example include asset returns, prices, or production decisions. The econometrician observes the data point r_t and aims to track the hidden state $s_t = (M_t, \Pi_t)$ of the learning economy.

We can apply the SOS filter to estimate the distribution of the state of the learning economy conditional on the observed data and the structural parameter θ . We propose an estimation procedure for θ based on indirect inference, a method introduced by Gouriéroux, Monfort and Renault (1993) and Smith (1993) that imputes the structural parameters of a model via an auxiliary estimator (e.g. Calzolari, Fiorentini and Sentana 2004; Czellar, Karolyi and Ronchetti 2007; Czellar and Ronchetti 2010; Dridi, Guay and Renault 2007; Genton and Ronchetti 2003; Heggland and Frigessi 2004). In our context, the full-information version of the economy, in which the state of nature M_t is directly observed by the agent, is a natural building block of the auxiliary estimator. When the state of nature takes finitely many values, the Bayesian filter and the likelihood of the full-information model are available analytically (e.g. Hamilton, 1989). Similarly, when

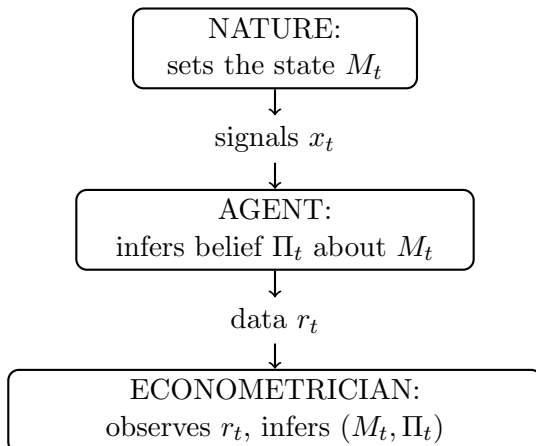


Figure 1: Information structure.

the state of nature M_t has an infinite support, a full-information economy with discretized M_t can be used. Given these properties, we define the auxiliary estimator by expanding the full-information economy’s maximum likelihood estimator with a set of statistics that the incomplete-information model is designed to capture.

We demonstrate the good performance of our techniques on a structural model of daily equity returns. Because the rich dynamics of the return series requires a large state space, we base our analysis on the multifrequency learning economy of Calvet and Fisher (“CF” 2007). We verify by Monte Carlo simulation that the SOS filter accurately tracks the state of the learning economy and provides remarkably precise estimates of the likelihood function. The indirect inference estimator is also shown to perform well in finite samples. We estimate the structural model on the daily excess returns of the CRSP U.S. value-weighted index between 1926 and 1999. For the out-of-sample period (2000-2009), the incomplete-information model provides accurate value-at-risk forecasts, which significantly outperform the predictions obtained from historical simulations, GARCH(1,1), and the full-information (FI) model.

The paper is organized as follows. Section 2 defines the SOS filter for general state space models. In section 3, we develop an indirect inference estimator for recursive learning economies. Section 4 applies these methods to a multifrequency investor learning model; we verify the accuracy of our inference methodology by Monte Carlo simulations, and conduct inference on the daily returns of a U.S. aggregate equity index between 1926 and 2009.

Section 5 concludes.

2 The State-Observation Sampling (SOS) Filter

2.1 Definition

We consider a discrete-time stochastic system defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Time is discrete and indexed by $t = 0, 1, \dots, \infty$. We consider a Markov process s_t defined on a measurable space $(\mathcal{S}, \mathcal{F}_S)$. For expositional simplicity, we assume in this subsection that $\mathcal{S} = \mathbb{R}^{n_S}$.

The econometrician receives every period an observation $r_t \in \mathbb{R}^{n_R}$. Let $R_{t-1} = (r_1, \dots, r_{t-1})$ denote the vector of observations up to date $t - 1$. The building block of our model is the conditional density of (s_t, r_t) given (s_{t-1}, R_{t-1}) :

$$f_{S,R}(s_t, r_t | s_{t-1}, R_{t-1}). \quad (2.1)$$

Let f_{S_0} denote a prior over the state space. The inference problem consists of estimating the density of the latent state s_t conditional on the set of current and past observations:

$$f_S(s_t | R_t)$$

at all $t \geq 1$.

A large literature proposes estimation by way of a particle filter, that is a finite set of points $(s_t^{(1)}, \dots, s_t^{(N)})$ that targets $f_S(s_t | R_t)$. The sampling importance resampling method of Gordon, Salmond, and Smith (1993) is based on Bayes' rule:

$$f_S(s_t | R_t) = \frac{f_R(r_t | s_t, R_{t-1}) f_S(s_t | R_{t-1})}{f_R(r_t | R_{t-1})}.$$

The recursive construction begins by drawing N independent states $s_0^{(1)}, \dots, s_0^{(N)}$ from f_{S_0} . Given the date- $(t - 1)$ filter $(s_{t-1}^{(1)}, \dots, s_{t-1}^{(N)})$, the construction of the date- t filter proceeds in two steps. First, we sample $\tilde{s}_t^{(n)}$ from $s_{t-1}^{(n)}$ using the transition kernel of the Markov process. Second, in the resampling step, we sample N particles $(s_t^{(1)}, \dots, s_t^{(N)})$

from $(\tilde{s}_t^{(n)}, \dots, \tilde{s}_t^{(N)})$ with normalized importance weights

$$p_t^{(n)} = \frac{f_R(r_t | \tilde{s}_t^{(n)}, R_{t-1})}{\sum_{n'=1}^N f_R(r_t | \tilde{s}_t^{(n')}, R_{t-1})}. \quad (2.2)$$

Under a wide range of conditions, the sample mean $N^{-1} \sum_{n=1}^N \Phi(s_t^{(n)})$ converges to $\mathbb{E}[\Phi(s_t) | R_t]$ for any bounded measurable function Φ .³

The sampling and importance resampling algorithm, and its various refinements, assume that the observation density $f_R(r_t | s_t, R_{t-1})$ is readily available up to a normalizing constant. This is a restrictive assumption in many applications, such as the incomplete-information economies considered in later sections.

We propose a solution to this difficulty when it is possible to simulate from (2.1). Our filter makes no assumption on the tractability of $f_{S,R}(\cdot | s_{t-1}, R_{t-1})$, and in fact does not even require that the transitional kernel of the Markov state s_t be available explicitly. The principle of our new filter is to simulate from each $s_{t-1}^{(n)}$ a state-observation pair $(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})$, and then select particles $\tilde{s}_t^{(n)}$ associated with pseudo-observations $\tilde{r}_t^{(n)}$ that are close to the actual data point r_t . The definition of the importance weights is based on Bayes' rule applied to the joint distribution of $\tilde{r}_t^{(n)}, \tilde{s}_t^{(n)}, s_{t-1}^{(n)}$ conditional on R_t :

$$\tilde{r}_t^{(n)}, \tilde{s}_t^{(n)}, s_{t-1}^{(n)} | R_t \sim \frac{\delta(r_t - \tilde{r}_t^{(n)}) f_{S,R}(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)} | s_{t-1}^{(n)}, R_{t-1}) f_S(s_{t-1}^{(n)} | R_{t-1})}{f_R(r_t | R_{t-1})}, \quad (2.3)$$

where δ denotes the Dirac distribution on \mathbb{R}^{n_R} . Since the Dirac distribution produces degenerate weights, we consider a kernel K with the following properties.

Assumption 1 (Kernel). *The function $K : \mathbb{R}^{n_R} \rightarrow \mathbb{R}_{++}$ satisfies:*

- (i) $\int K(u) du = 1$;
- (ii) $\int u K(u) du = 0$;
- (iii) $A(K) = \int \|u\|^2 K(u) du < \infty$;
- (iv) $B(K) = \int [K(u)]^2 du < \infty$.

³See Crisan and Doucet (2002) for an excellent survey on the convergence of particle filters.

For any $r \in \mathbb{R}^{n_R}$, let

$$K_{h_t}(r) = \frac{1}{h_t^{n_R}} K\left(\frac{r}{h_t}\right)$$

denote the corresponding kernel with bandwidth h_t at date t . The kernel K_{h_t} converges to the Dirac distribution as h_t goes to zero, which we use to approximate (2.3). This suggests the following algorithm.

SOS filter

Step 1 (State-observation sampling): For every $n = 1, \dots, N$, we simulate a state-observation pair $(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})$ from $f_{S,R}(\cdot | s_{t-1}^{(n)}, R_{t-1})$.

Step 2 (Importance weights): We observe the new data point r_t and compute

$$p_t^{(n)} = \frac{K_{h_t}(r_t - \tilde{r}_t^{(n)})}{\sum_{n'=1}^N K_{h_t}(r_t - \tilde{r}_t^{(n')})}, \quad n = 1, \dots, N.$$

Step 3 (Multinomial resampling): For every $n = 1, \dots, N$, we draw $s_t^{(n)}$ from $\tilde{s}_t^{(1)}, \dots, \tilde{s}_t^{(N)}$ with importance weights $p_t^{(1)}, \dots, p_t^{(N)}$.

The state-observation pairs $\{(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})\}_{n=1, \dots, N}$ constructed in step 1 provide a discrete approximation to the conditional distribution of (s_t, r_t) given the data R_{t-1} . In step 2, we construct a measure of the proximity between the pseudo and the actual data points, and in Step 3 we select particles for which this measure is large. The variance of multinomial resampling in step 3 can be reduced and computational speed can be improved by alternatives such as residual (Liu and Chen, 1998) or stratified (Kitagawa, 1996) resampling. In section 4, we obtain good results with a combined residual-stratified approach.⁴ The convergence proof below applies equally well to these alternatives.

⁴We select $\sum_{n=1}^N \lfloor N p_t^{(n)} \rfloor$ particles deterministically by setting $\lfloor N p_t^{(n)} \rfloor$ particles equal to $\tilde{s}_t^{(n)}$ for every $n \in \{1, \dots, N\}$, where $\lfloor \cdot \rfloor$ denotes the floor of a real number. The remaining $N_{r,t} = N - \sum_{n=1}^N \lfloor N p_t^{(n)} \rfloor$ particles are selected by the stratified sampling that produces $\tilde{s}_t^{(n)}$ with probability $q_t^{(n)} = (N p_t^{(n)} - \lfloor N p_t^{(n)} \rfloor) / N_{r,t}$, $n = 1, \dots, N$. That is, for every $k \in \{1, \dots, N_{r,t}\}$, we draw \tilde{U}_k from the uniform distribution on $(\frac{k-1}{N_{r,t}}, \frac{k}{N_{r,t}}]$, and select the particle $\tilde{s}_t^{(n)}$ such that $\tilde{U}_k \in (\sum_{j=1}^{n-1} q_t^{(j)}, \sum_{j=1}^n q_t^{(j)}]$.

2.2 Extension and Convergence

The SOS filter easily extends to the case of a general measurable state space \mathcal{S} . The building blocks of the model are the conditional probability measure of (s_t, r_t) given (s_{t-1}, R_{t-1}) :

$$g(\cdot | s_{t-1}, R_{t-1}),$$

and a prior measure λ_0 over the state space. The SOS filter targets the probability measure of the latent state s_t conditional on the set of current and past observations, $\lambda(\cdot | R_t)$. The SOS filter is defined as in Section 2, where in step 1 we sample $(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})$ from the conditional measure $g(\cdot | s_{t-1}^{(n)}, R_{t-1})$.

We now specify conditions under which for an arbitrary state space \mathcal{S} and a fixed history $R_T = (r_1, \dots, r_T)$, $T \leq \infty$, the SOS filter converges in mean squared error to the target $\lambda(\cdot | R_t)$ as the number of particles N goes to infinity.

Assumption 2 (Conditional Distributions). *The observation process satisfies the following hypotheses:*

(i) *the conditional density $f_R(\tilde{r}_t | s_{t-1}, R_{t-1})$ exists and*

$$\kappa_t = \sup\{f_R(\tilde{r}_t | s_{t-1}, R_{t-1}); (s_{t-1}, \tilde{r}_t) \in \mathcal{S} \times \mathbb{R}^{n_R}\} < \infty;$$

(ii) *the observation density $f_R(\tilde{r}_t | s_t, R_{t-1})$ is well-defined and there exists $\kappa'_t \in \mathbb{R}_+$ such that:*

$$|f_R(\tilde{r}_t | s_t, R_{t-1}) - f_R(r_t | s_t, R_{t-1}) - \frac{\partial f_R}{\partial r'_t}(r_t | s_t, R_{t-1})(\tilde{r}_t - r_t)| \leq \kappa'_t \|\tilde{r}_t - r_t\|^2$$

for all $(s_t, \tilde{r}_t) \in \mathcal{S} \times \mathbb{R}^{n_R}$ and $t \leq T$.

Assumption 3 (Bandwidth). *The bandwidth is a function of N , $h_t = h_t(N)$, and satisfies*

(i) $\lim_{N \rightarrow \infty} h_t(N) = 0$,

(ii) $\lim_{N \rightarrow \infty} N[h_t(N)]^{n_R} = +\infty$,

for all $t = 1, \dots, T$.

We establish the following result in the appendix.

Theorem 4 (Convergence of the SOS Filter). *Under assumptions 1 and 2 and for every t and $N \geq 1$, there exists $U_t(N) \in \mathbb{R}_+$ such that*

$$\mathbb{E} \left\{ \left[\frac{1}{N} \sum_{n=1}^N K_{h_t}(r_t - \tilde{r}_t^{(n)}) - f_R(r_t | R_{t-1}) \right]^2 \right\} \leq \frac{[f_R(r_t | R_{t-1})]^2}{4} U_t(N), \quad (2.4)$$

where the expectation is over all the realizations of the random particle method. Furthermore, for any bounded measurable function, $\Phi : \mathcal{S} \rightarrow \mathbb{R}$,

$$MSE_t = \mathbb{E} \left\{ \left[\frac{1}{N} \sum_{n=1}^N \Phi(s_t^{(n)}) - \mathbb{E}[\Phi(s_t) | R_t] \right]^2 \right\} \leq U_t(N) \|\Phi\|^2, \quad (2.5)$$

where $\|\Phi\| = \sup_{s \in \mathcal{S}} |\Phi(s)|$. If assumption 3 also holds, then

$$\lim_{N \rightarrow \infty} U_t(N) = 0,$$

and the filter converges in mean squared error. Furthermore, if the bandwidth sequence is of the form $h_t(N) = h_t(1)N^{-1/(n_R+4)}$, then $U_t(N)$ decays at rate $N^{-4/(n_R+4)}$ and the root mean squared error $MSE_t^{1/2}$ at rate $N^{-2/(n_R+4)}$ for all t .

By (2.4), the kernel estimator

$$\hat{f}_R(r_t | R_{t-1}) = \frac{1}{N} \sum_{n=1}^N K_{h_t}(r_t - \tilde{r}_t^{(n)}), \quad (2.6)$$

converges to the conditional density of r_t given past observations. Consequently, we can estimate the log-likelihood function by $\sum_{t=1}^T \ln \hat{f}_R(r_t | R_{t-1})$, and provide a plug-in bandwidth in the online Appendix. We will illustrate in section 4 the finite-sample accuracy of the SOS filter.

3 Recursive Learning Economies

We consider a class of discrete-time stochastic economies defined at $t = 0, \dots, \infty$ on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and parameterized by $\theta \in \Theta \subseteq \mathbb{R}^p$, $p \geq 1$.

3.1 Information Structure

In every period t , we define three levels of information, which respectively correspond to nature, a Bayesian agent, and the econometrician. Figure 1 illustrates the information structure.

3.1.1 Nature

A *state of nature* M_t drives the fundamentals of the economy. We assume that M_t follows a first-order Markov chain on the set of mutually distinct states $\{m^1(\theta), \dots, m^d(\theta)\}$. For every $i, j \in \{1, \dots, d\}$, we denote by $a_{i,j}(\theta) = \mathbb{P}(M_t = m^j(\theta) | M_{t-1} = m^i(\theta); \theta)$ the transition probability from state i to state j . We assume that the Markov chain M_t is irreducible, aperiodic, positive recurrent, and therefore ergodic. For notational simplicity, we henceforth drop the argument θ from the states m^j and transition probabilities $a_{i,j}$.

3.1.2 Agent

At the beginning of every period t , the agent observes a signal vector $x_t \in \mathbb{R}^{n_x}$, which is partially revealing on the state of nature M_t . The probability density function of the signal conditional on the state of nature, $f_X(x_t | M_t; \theta)$, is known to the agent. Let $X_t = (x_1, \dots, x_t)$ denote the vector of signals received by the agent up to date t . For tractability reasons, we make the following hypotheses.

Assumption 5 (Signal). *The signal satisfies the following conditions:*

- (a) $\mathbb{P}(M_t = m^j | M_{t-1} = m^i, X_{t-1}; \theta) = a_{i,j}$ for all i, j ;
- (b) $f_X(x_t | M_t, M_{t-1}, \dots, M_0, X_{t-1}; \theta) = f_X(x_t | M_t; \theta)$.

The agent knows the structural parameter θ , is Bayesian and uses X_t to compute the conditional probability of the states of nature.

Proposition 6 (Agent Belief). *The conditional probabilities $\Pi_t^j = \mathbb{P}(M_t = m^j | X_t; \theta)$ satisfy the recursion:*

$$\Pi_t^j = \frac{\omega^j(\Pi_{t-1}, x_t; \theta)}{\sum_{i=1}^d \omega^i(\Pi_{t-1}, x_t; \theta)} \text{ for all } j \in \{1, \dots, d\} \text{ and } t \geq 1, \quad (3.1)$$

where $\Pi_{t-1} = (\Pi_{t-1}^1, \dots, \Pi_{t-1}^d)$ and $\omega^j(\Pi_{t-1}, x_t; \theta) = f_X(x_t | M_t = m^j; \theta) \sum_{i=1}^d a_{i,j} \Pi_{t-1}^i$.

In applications, the agent values assets or makes financial, production or purchasing decisions as a function of the belief vector Π_t . Our methodology easily extends to learning models with non-Bayesian agents, as in Brandt, Zeng, and Zhang (2004) and Cecchetti Lam and Mark (2000).

The *state of the learning economy* at a given date t is the mixed variable $s_t = (M_t, \Pi_t)$. The state space is therefore

$$\mathcal{S} = \{m^1, \dots, m^d\} \times \Delta_+^{d-1}, \quad (3.2)$$

where $\Delta_+^{d-1} = \{\Pi \in \mathbb{R}_+^d | \sum_{i=1}^d \Pi_i = 1\}$ denotes the $(d-1)$ -dimensional unit simplex.

Proposition 7 (State of the Learning Economy). *The state of the learning economy, $s_t = (M_t, \Pi_t)$, is first-order Markov. It is ergodic if the transition probabilities between states of nature are strictly positive: $a_{i,j} > 0$ for all i, j , and the signal's conditional probability density functions $f_X(x | M_t = m^j; \theta)$ are strictly positive for all $x \in \mathbb{R}^{n_X}$ and $j \in \{1, \dots, d\}$.*

The state of the learning economy s_t preserves the first-order Markov structure of the state of nature M_t . By Bayes'rule (3.1), the transition kernel of the Markov state s_t is sparse when the dimension of the signal, n_X , is lower than the number of states of nature: $n_X < d$. The state s_t is nonetheless ergodic for all values n_X and d under the conditions stated in Proposition 7, which guarantees that the economy is asymptotically independent of the initial state s_0 .

3.1.3 Econometrician

Each period, the econometrician observes a data point $r_t \in \mathbb{R}^{n_R}$, which is assumed to be a deterministic function of the agent's signal and conditional probabilities over states of

nature:

$$r_t = \mathcal{R}(x_t, \Pi_t, \Pi_{t-1}; \theta). \quad (3.3)$$

We include Π_{t-1} in this definition to accommodate the possibility that r_t is a growth rate or return. The parameter vector $\theta \in \mathbb{R}^p$ specifies the states of nature m^1, \dots, m^d , their transition probabilities $(a_{i,j})_{1 \leq i, j \leq d}$, the signal's conditional density $f_X(\cdot | M_t, \theta)$, and the data function $\mathcal{R}(x_t, \Pi_t, \Pi_{t-1}; \theta)$. In some applications, it may be useful to add measurement error in (3.3); the estimation procedure of the next section applies equally well to this extension.

3.2 Estimation

We assume that the data $R_T = (r_1, \dots, r_T)$ is generated by the incomplete-information (II) economy with parameter θ^* described above. Estimation faces several challenges. The transition kernel of the Markov state s_t and the log-likelihood function $\mathcal{L}_{II}(\theta | R_T)$ are not available analytically. Furthermore, the observation density $f_R(r_t | s_t, R_{t-1})$ is not available in closed form either because the signal x_t , drives the data point $r_t = \mathcal{R}(x_t, \Pi_t, \Pi_{t-1}; \theta)$ both directly and indirectly through the belief $\Pi_t = \Pi_t(x_t, \Pi_{t-1})$, creating a highly nonlinear relationship between the state and the observation.

The learning model can, however, be conveniently simulated. Given a state $s_{t-1} = (M_{t-1}, \Pi_{t-1})$, we can: (i) sample M_t from M_{t-1} using the transition probabilities $a_{i,j}$; (ii) sample the signal x_t from $f_X(\cdot | M_t; \theta)$; (iii) apply Bayes'rule (3.1) to impute the agent's belief Π_t ; and (iv) compute the simulated data point $\tilde{r}_t = \mathcal{R}(x_t, \Pi_t, \Pi_{t-1}; \theta)$. Estimation can therefore proceed by simulation-based methods. Simulated ML based on the SOS filter is a possible approach. As we will see in section 4, however, an accurate approximation of the log-likelihood value $\hat{\mathcal{L}}_{II}(\theta | R_T) = \sum_{t=1}^T \ln[N^{-1} \sum_{n=1}^N K_{h_t}(r_t - \tilde{r}_t^{(n)})]$ may require a large number of particles. For situations where simulated ML is too computational⁵, we now propose an alternative approach based on indirect inference.

For each learning model $\theta \in \Theta$, we can define an auxiliary *full information* (FI) model in which the agent observes both the state of nature M_t and the signal x_t . Her conditional probabilities are then $\Pi_t^j = \mathbb{P}(M_t = m^j | X_t, M_t; \theta)$ for all j . The belief vector reduces to $\Pi_t = \mathbf{1}_{M_t}$, where $\mathbf{1}_{M_t}$ denotes the vector whose j^{th} component is equal to 1

⁵For instance in the empirical example considered in section 4, we use an SOS filter of size $N = 10^7$ and a dataset of about 20,000 observations. One evaluation of the likelihood function requires the evaluation 200 billion kernels $K_{h_t}(\cdot)$. Since a typical optimization requires about 500 function evaluations, the simulated ML estimation of the II model would require the evaluation of 100 trillion kernels.

if $M_t = m^j$ and 0 otherwise, and by (3.3) the full information data point is defined by $r_t = \mathcal{R}(x_t, \mathbb{1}_{M_t}, \mathbb{1}_{M_{t-1}}; \theta)$. The FI model can have less parameters than the II model because of the simplification in Π_t . We therefore consider that the auxiliary FI model is parameterized by $\phi \in \mathbb{R}^q$, where $1 \leq q \leq p$.

Assumption 8 (Auxiliary Full-Information Economies). *The probability density functions $f_{i,j}(r_t; \phi) = f_{R,FI}(r_t | M_t = m^j, M_{t-1} = m^i, \phi)$ are available analytically for all $i, j \in \{1, \dots, d\}$.*

Proposition 9 (Full-Information Likelihood). *Under assumption 8, the log-likelihood function $\mathcal{L}_{FI}(\phi | R_T)$ is available analytically.*

The ML estimator of the full-information economy

$$\hat{\phi}_T = \arg \max_{\phi} \mathcal{L}_{FI}(\phi | R_T) \in \mathbb{R}^q$$

can therefore be conveniently computed.

The indirect-inference estimation of the structural learning model proceeds in two steps. First, we define an auxiliary estimator that includes the full-information MLE. If $q < p$, we also consider a set of $p - q$ statistics $\hat{\eta}_T$ that quantify features of the dataset R_T that the learning model is designed to capture. The *auxiliary estimator* is defined by

$$\hat{\mu}_T = \begin{bmatrix} \hat{\phi}_T \\ \hat{\eta}_T \end{bmatrix} \in \mathbb{R}^p. \quad (3.4)$$

By construction, $\hat{\mu}_T$ contains as many parameters as the structural parameter θ .⁶

Second, for any admissible parameter θ , we can simulate a sample path $R_{ST}(\theta)$ of length ST , $S \geq 1$, and compute the corresponding pseudo-auxiliary estimator:

$$\hat{\mu}_{ST}(\theta) = \begin{bmatrix} \hat{\phi}_{ST}(\theta) \\ \hat{\eta}_{ST}(\theta) \end{bmatrix}, \quad (3.5)$$

⁶We focus on the exactly identified case to simplify the exposition and because earlier evidence indicates that parsimonious auxiliary models tend to provide more accurate inference in finite samples (e.g. Andersen, Chung, and Sorensen, 1999; Czellar and Ronchetti, 2010). Our approach naturally extends to the overidentified case, which may be useful in cases where it is economically important to match a larger set of statistics.

where $\hat{\phi}_{ST}(\theta) = \arg \max_{\phi} \mathcal{L}_{FI}[\phi | R_{ST}(\theta)]$. We define the indirect inference estimator $\hat{\theta}_T$ by:

$$\hat{\theta}_T = \arg \min_{\theta} [\hat{\mu}_{ST}(\theta) - \hat{\mu}_T]' \Omega [\hat{\mu}_{ST}(\theta) - \hat{\mu}_T], \quad (3.6)$$

where Ω is a positive definite weighting matrix. When the calculation of the full-information MLE is expensive, the numerical implementation can be accelerated by the efficient method of moments, as is discussed in the appendix.

Our methodology builds on the fact that the full-information economy can be efficiently estimated by ML and is therefore a natural candidate auxiliary model. Moreover, the theoretical investigation of a learning model often begins with the characterization of the FI case, so the estimation method we are proposing follows the natural progression commonly used in the literature.

We assume that the assumptions 10–12 given in the appendix hold. Gouriéroux et al. (1993) and Gouriéroux and Monfort (1996) show that under these conditions and assuming the structural model θ^* , the auxiliary estimator $\hat{\mu}_T$ converges in probability to a deterministic function $\mu(\theta^*)$, called the *binding function*, and $\sqrt{T} [\hat{\mu}_T - \mu(\theta^*)] \xrightarrow{d} \mathcal{N}(0, W^*)$, where W^* is defined in the appendix. Furthermore, when S is fixed and T goes to infinity, the estimator $\hat{\theta}_T$ is consistent and asymptotically normal:

$$\sqrt{T}(\hat{\theta}_T - \theta^*) \xrightarrow{d} \mathcal{N}(0, \Sigma),$$

where

$$\Sigma = \left(1 + \frac{1}{S}\right) \left[\frac{\partial \mu(\theta^*)}{\partial \theta'}\right]^{-1} W^* \left[\frac{\partial \mu(\theta^*)'}{\partial \theta}\right]^{-1}. \quad (3.7)$$

The appendix further discusses the numerical implementation of this method.

In this section, we have assumed that the state of nature takes finitely many values. When M_t has an infinite support, we can discretize its distribution and use the corresponding full-information discretized economy as an auxiliary model. The definition and properties of the indirect inference estimator are otherwise identical.

4 Inference in an Asset Pricing Model with Investor Learning

We now apply our methodology to a consumption-based asset pricing model. We adopt the Lucas tree economy with regime-switching fundamentals of CF (2007), which we use to specify the dynamics of daily equity returns.

4.1 Specification

4.1.1 Dynamics of the State of Nature

The rich dynamics of daily returns requires a large state space. For this reason, we consider that the state is a vector containing \bar{k} components:

$$M_t = (M_{1,t}, \dots, M_{\bar{k},t})' \in \mathbb{R}_+^{\bar{k}},$$

which follows a binomial Markov Switching Multifractal (CF 2001, 2004, 2008). The components are mutually independent across k . Let M denote a Bernoulli distribution that takes either a high value m_0 or a low value $2 - m_0$ with equal probability. Given a value $M_{k,t}$ for the k^{th} component at date t , the next-period multiplier $M_{k,t+1}$ is either:

$$\begin{cases} \text{drawn from the distribution } M \text{ with probability } \gamma_k, \\ \text{equal to its current value } M_{k,t} \text{ with probability } 1 - \gamma_k. \end{cases}$$

Since each component of the state vector can take two possible values, the state space contains $d = 2^{\bar{k}}$ elements m^1, \dots, m^d . The transition probabilities γ_k are parameterized by

$$\gamma_k = 1 - (1 - \gamma_{\bar{k}})^{b^{k-\bar{k}}}, \quad k = 1, \dots, \bar{k},$$

where $b > 1$. Thus, $\gamma_{\bar{k}}$ controls the persistence of the highest-frequency component and b determines the spacing between frequencies.

4.1.2 Bayesian Agent

The agent receives an exogenous consumption stream $\{C_t\}$ and prices the stock, which is a claim on an exogenous dividend stream $\{D_t\}$. Every period, the agent observes a signal

$x_t \in \mathbb{R}^{\bar{k}+2}$ consisting of dividend growth:

$$x_{1,t} = \ln(D_t/D_{t-1}) = g_D - \frac{\sigma_D^2(M_t)}{2} + \sigma_D(M_t)\varepsilon_{D,t}, \quad (4.1)$$

consumption growth:

$$x_{2,t} = \ln(C_t/C_{t-1}) = g_C + \sigma_C\varepsilon_{C,t}, \quad (4.2)$$

and a noisy version of the state:

$$x_{i+2,t} = M_{i,t} + \sigma_\delta z_{i,t}, \quad i = 1, \dots, \bar{k}. \quad (4.3)$$

The noise parameter $\sigma_\delta \in \mathbb{R}_+$ controls information quality. The stochastic volatility of dividends is given by:

$$\sigma_D(M_t) = \bar{\sigma}_D \left(\prod_{k=1}^{\bar{k}} M_{k,t} \right)^{1/2}, \quad (4.4)$$

where $\bar{\sigma}_D \in \mathbb{R}_+$. The innovations $\varepsilon_{C,t}$, $\varepsilon_{D,t}$, and z_t are jointly normal and have zero means and unit variances. We assume that $\varepsilon_{C,t}$ and $\varepsilon_{D,t}$ have correlation $\rho_{C,D}$, and that all the other correlation coefficients are zero.

Learning about the volatility state M_t is an asymmetric process. For expositional simplicity, assume that the noise parameter σ_δ is large, so that investors learn about M_t primarily through the dividend growth. Because large realizations of dividend are implausible in a low-volatility regime, learning about a volatility increase tends to be abrupt. Conversely, when volatility switches from a high to a low state, the agent learns only gradually that volatility has gone down because realizations of dividend growth near the mean are likely outcomes under any M_t .

The agent has isoelastic expected utility, $U_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \delta^t C_t^{1-\alpha}/(1-\alpha)$, where δ is the discount rate and α is the coefficient of relative risk aversion. In equilibrium, the log interest rate is constant. The stock's price-dividend ratio is negatively related to volatility and linear in the belief vector:

$$Q(\Pi_t) = \sum_{j=1}^d Q(m^j)\Pi_t^j. \quad (4.5)$$




Figure 2: Learning Model Simulation. This figure illustrates a sample path simulated from the multifrequency learning model. Each panel corresponds to a different level of information. Nature's price-dividend ratio $Q(M_t)$ is plotted in the top panel, the agent's price-dividend ratio $Q(\Pi_t)$ in the middle panel, and the return r_t (computed by the agent and observed by the econometrician) in the bottom panel.

where the linear coefficients $Q(m^j)$ are available analytically.⁷

4.1.3 Econometric Specification of Stock Returns

The econometrician observes the log excess return process:

$$r_t = \ln \left[\frac{1 + Q(\Pi_t)}{Q(\Pi_{t-1})} \right] + x_{1,t} - r_f. \quad (4.6)$$

Since learning about volatility is asymmetric, the stock price falls abruptly following a volatility increase (bad news), but will increase only gradually after a volatility decrease (good news). The noise parameter σ_δ therefore controls the skewness of stock returns.

4.2 Accuracy of the SOS filter

We now present the results of Monte Carlo simulations of the SOS particle filter. To simplify the exposition, we consider one-dimensional aggregates of M_t and Π_t , which summarize economically meaningful information. Specifically, if the agent knew the true state of nature, she would set the price-dividend ratio equal to $Q(M_t) = Q(m^j)$ if $M_t = m^j$, as implied by (4.5); we therefore call $Q(M_t)$ *nature's P/D ratio*. By contrast, the market $Q(\Pi_t)$ aggregates the agent's beliefs in the incomplete-information model; for this reason, we refer to it as the *agent's price-dividend ratio*.

We generate a sample of size $T = 20,000$ periods from the learning model (4.6) with $\bar{k} = 3$ volatility components and fixed parameter values.⁸ Figure 2 illustrates the last 1,000 periods of the simulated sample. We report nature's price-dividend ratio in the top panel, the agent's price-dividend ratio in the middle panel, and the return (computed by the agent and observed by the econometrician) in the bottom panel.

⁷The price-dividend ratio is given by

$$\sum_{n=1}^{\infty} \delta^n \mathbb{E} \left[\left(\frac{C_{t+n}}{C_t} \right)^{-\alpha} \frac{D_{t+n}}{D_t} \middle| X_t \right] = \sum_{n=1}^{\infty} \mathbb{E} \left[\prod_{h=1}^n e^{g_D - r_f - \alpha \rho_{C,D} \sigma_C \sigma_D (M_{t+h})} \middle| X_t \right],$$

where $r_f = -\ln(\delta) + \alpha g_C - \alpha^2 \sigma_C^2 / 2$ is the log interest rate. Since volatility is persistent, a high level of volatility at date t implies high forecasts of future volatility, and therefore a low period- t price-dividend ratio. The linear coefficients are given by $(Q(m^1), \dots, Q(m^d))' = (I - B)^{-1} \iota - \iota$, where $B = (b_{ij})_{1 \leq i, j \leq d}$ is the matrix with components $b_{ij} = a_{i,j} \exp [g_D - r_f - \alpha \rho_{C,D} \sigma_C \sigma_D (m^j)]$ and $\iota = (1, \dots, 1)'$.

⁸Specifically, we set $m_0 = 1.7$, $\gamma_{\bar{k}} = 0.06$, $b = 2$ and $\sigma_\delta = 1$, the consumption drift to $g_C = 0.75$ basis point (bp) (or 1.18% per year), log interest rate to $r_f = 0.42$ bp per day (1% per year), excess dividend growth equal to $g_D - r_f = 0.5$ bp per day (about 1.2% per year), consumption volatility to $\sigma_C = 0.189\%$ (or 2.93% per year), and dividend volatility $\bar{\sigma}_D = 0.70\%$ per day (about 11% per year). The correlation coefficient is set equal to $\rho_{C,D} = 0.6$, and α is chosen such that the mean of the linear coefficients in (4.5) satisfy $\bar{Q} = d^{-1} \sum_{i=1}^d Q(m^i) = 6000$ in daily units (25 in yearly units).




Figure 3: Accuracy of the SOS Filter. This figure illustrates the estimated log-likelihood function (left panel) and the efficiency measures $R_{Q(\Pi)}^2$ and $R_{Q(M)}^2$ (right panel) as a function of the filter size N .

We apply to the entire simulated sample the SOS filter with the quasi-Cauchy kernel and bandwidth derived in the online Appendix. The left panel of Figure 3 illustrates the estimated log-likelihood as a function of the filter size N . In the right panel, we report the pseudo R^2 :

$$R_{Q(\Pi)}^2 = 1 - \frac{\sum_{t=1}^T [\hat{Q}(\Pi_t) - Q(\Pi_t)]^2}{\sum_{t=1}^T [\hat{Q}(\Pi_t) - \bar{Q}(\Pi)]^2},$$

where $\hat{Q}(\Pi_t) = \sum_{n=1}^N Q(\Pi_t^{(n)})/N$ and $\bar{Q}(\Pi) = \sum_{t=1}^T Q(\Pi_t)/T$. We similarly compute $R_{Q(M)}^2$ for nature's price-dividend ratios using $\{Q(M_t^{(n)})\}$. The figure shows that both the estimated log-likelihood and the coefficients of determination increase with the filter size N and settle down for $N \geq 10^6$. The coefficient of determination reaches 67.6% for $Q(M)$ and 71.5% for $Q(\Pi)$. Thus, the agent's P/D ratio is better estimated than nature's P/D ratio, as the information structure in Figure 1 suggests.

The true value of the likelihood function is unknown for the example considered in Figure 3. For this reason, we now consider the full-information version of the model, which, by Proposition 9, has a closed-form likelihood. We generate from the full-information model a sample of $T = 20,000$ periods. The analytical expression of the log-likelihood implies that $\mathcal{L}_{FI} = 79,691.5$. In the right column of Table 2, we report the sample mean and the root mean squared error of fifty log-likelihood estimates computed using SOS. The relative estimation error $\text{RMSE}/\mathcal{L}_{FI}$ is 0.024%, 0.006% and 0.002% when using, respectively, $N = 10^5, 10^6$ and 10^7 particles. The estimates of the FI log-likelihood obtained using SOS are therefore remarkably precise.

We now verify that the SOS filter defeats the curse of dimensionality with respect to the size of the state space. Table 1 reports the topological dimension of the state space, $\dim \mathcal{S}$, under incomplete and full information. By construction, the log-likelihood function satisfies the continuity property: $\lim_{\sigma_\delta \rightarrow 0} \mathcal{L}_{II}(m_0, \gamma_{\bar{k}}, b, \sigma_\delta | R_T) = \mathcal{L}_{FI}(m_0, \gamma_{\bar{k}}, b | R_T)$. The first three columns in Table 2 report summary statistics of log-likelihood estimates of \mathcal{L}_{II} obtained for $\sigma_\delta \in \{1, 0.1, 0.01\}$. The accuracy of SOS is nearly identical for the full-information model and for the learning model with $\sigma_\delta = 0.01$. With $N = 10^7$ particles, the RMSE of the SOS filter is even slightly smaller for the II specification $\sigma_\delta = 0.01$ than for the full-information model, even though II has a much larger state space. These findings confirm the result of Theorem 4 that the convergence rate of SOS is independent of the dimension of the state space.

Table 1: DIMENSION OF THE STATE SPACE^a

| | Incomplete Information | Full Information |
|------------------------------|---|-----------------------|
| State space \mathcal{S} | $\{m^1, \dots, m^d\} \times \Delta_+^{d-1}$ | $\{m^1, \dots, m^d\}$ |
| Dimension $\dim \mathcal{S}$ | $d - 1$ | 0 |

^aThis table reports the topological dimension of the state space under full and incomplete information. In the multifrequency volatility case, we know that $d = 2^{\bar{k}}$, where \bar{k} denote the number of volatility frequencies.

Table 2: PRECISION OF THE SOS LOG-LIKELIHOOD ESTIMATES^a

| | II ($\dim \mathcal{S} = 7$) | | | FI ($\dim \mathcal{S} = 0$) |
|------------------|-------------------------------|-----------------------|------------------------|-------------------------------|
| | $\sigma_\delta = 1$ | $\sigma_\delta = 0.1$ | $\sigma_\delta = 0.01$ | |
| Mean, $N = 10^5$ | 79,514.1 | 79,674.0 | 79,673.1 | 79,673.4 |
| Mean, $N = 10^6$ | 79,523.2 | 79,686.3 | 79,687.4 | 79,687.3 |
| Mean, $N = 10^7$ | 79,525.1 | 79,690.8 | 79,690.9 | 79,690.4 |
| RMSE, $N = 10^5$ | 177.6 | 18.3 | 19.5 | 18.9 |
| RMSE, $N = 10^6$ | 168.4 | 6.3 | 5.2 | 4.9 |
| RMSE, $N = 10^7$ | 166.4 | 1.1 | 1.3 | 1.7 |

^aWe report summary statistics for 50 simulated log-likelihoods estimated on a fixed sample path of $T = 20,000$ periods from the FI model. The true log-likelihood is $\mathcal{L}_{FI} = 79,691.5$. The simulated log-likelihoods are based on an SOS filter and a learning model with $\sigma_\delta \in \{0.01, 0.1, 1\}$.

4.3 Indirect Inference Estimator

We now develop an estimator for the vector of structural parameters:

$$\theta = (m_0, \gamma_{\bar{k}}, b, \sigma_\delta)' \in [1, 2] \times (0, 1) \times [1, \infty) \times \mathbb{R}_+,$$

where m_0 controls the variability of dividend volatility, $\gamma_{\bar{k}}$ the transition probability of the most transitory volatility component, b the spacing of the transition probabilities, and σ_δ the precision of the signal received by the representative agent. As is traditional in the asset pricing literature, we calibrate all the other parameters on aggregate consumption data and constrain the mean price-dividend ratio to a plausible long-run value

$$\mathbb{E}[Q(\Pi_t)] = \bar{Q}, \tag{4.7}$$

where \bar{Q} is set equal to 25 in yearly units.⁹

The learning economy is specified by $p = 4$ parameters, $\theta = (m_0, \gamma_{\bar{k}}, b, \sigma_\delta)'$, while the FI economy is specified by $q = 3$ parameters, $\phi = (m_0, \gamma_{\bar{k}}, b)'$. For this reason, the definition of the auxiliary estimator requires an additional statistic $\hat{\eta}_T \in \mathbb{R}$. Since the noise parameter σ_δ controls the skewness of excess returns, the third moment seems like a natural choice. We are concerned, however, that the third moment may be too noisy to produce an efficient estimator of θ . For this reason, we consider an alternative based on the observation that by restriction (4.7), the mean return is nearly independent of the structural parameter:

$$\mathbb{E}(r_t) \approx \ln(1 + 1/\bar{Q}) + g_D - r_f - \bar{\sigma}_D^2, \tag{4.8}$$

as is verified in the online appendix. Since the mean is fixed, the median can be used as a robust measure of skewness. The auxiliary estimator $\hat{\mu}_T = (\hat{\phi}_T, \hat{\eta}_T)'$ is defined by expanding the ML estimator of the full-information economy, $\hat{\phi}_T$, with either the third moment ($\hat{\eta}_T = T^{-1} \sum_{t=1}^T r_t^3$) or median ($\hat{\eta}_T = \text{median}\{r_t\}$) of returns.

In Figure 4, we illustrate the relation between the median-based auxiliary estimator $\hat{\mu}_T$ and the structural parameter θ on a long simulated sample of length $ST = 10^7$. The graphs

⁹The calibrated parameters are the same as in the previous subsection. An alternative approach would be to estimate all the parameters of the learning economy on aggregate excess return data. In the 2005 NBER version of their paper, CF applied this method to the FI model and obtained broadly similar results to the ones reported in the published version. This alternative approach has the disadvantage of not taking into account the economic constraints imposed by the model, and we do not pursue it here.

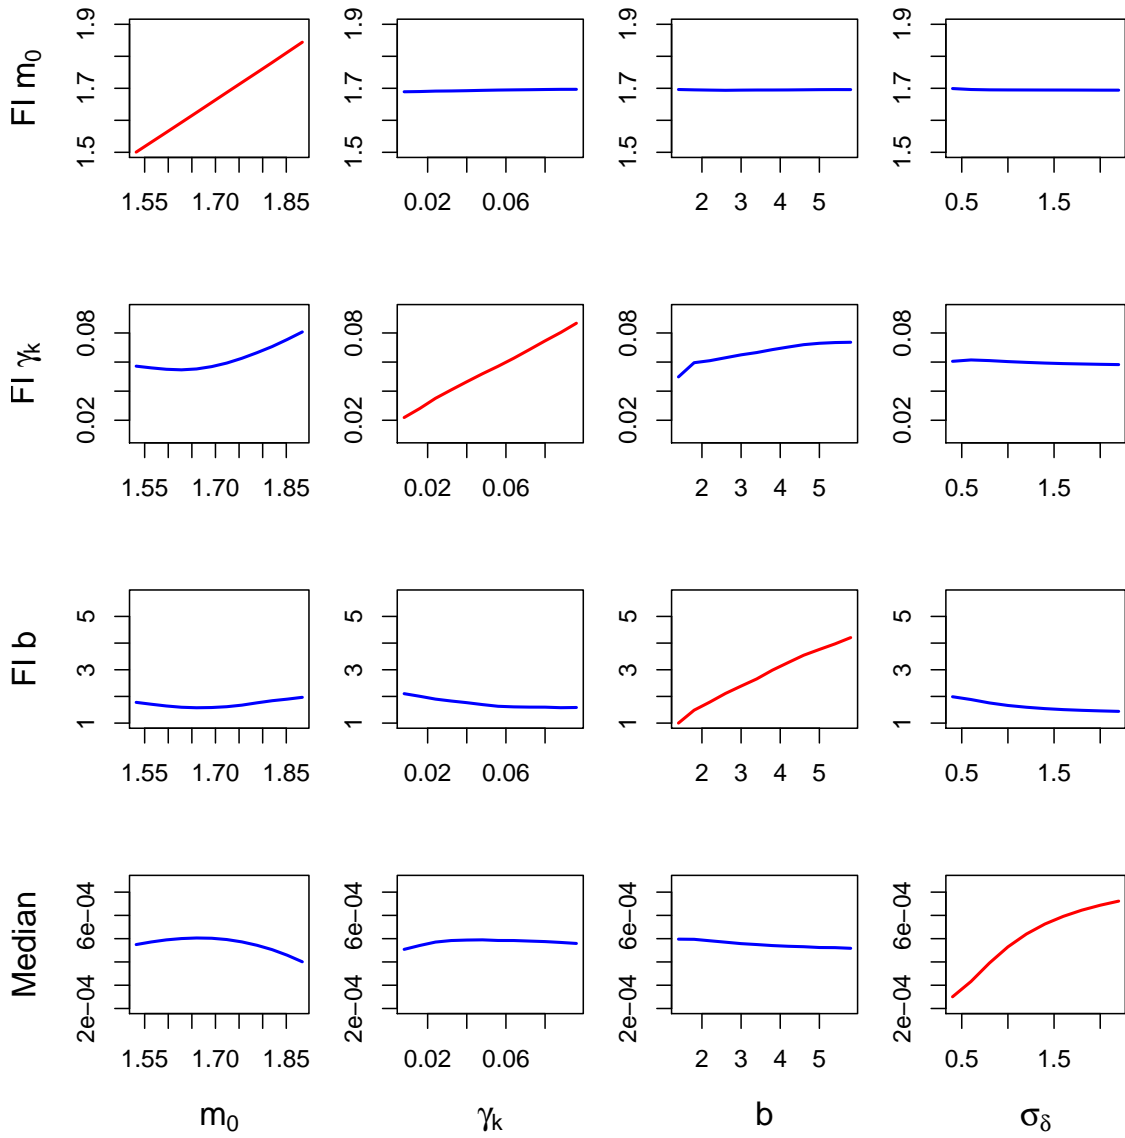


Figure 4: Auxiliary Estimator. This figure illustrates the relation between the median-based auxiliary estimator and the structural parameter θ . In each column, one structural parameter is allowed to vary while the other three parameters are set to their reference values. The auxiliary estimate reported for every θ is obtained from a simulated sample of length 10^7 generated under the learning model θ with $\bar{k} = 3$ volatility components.

can be viewed as cuts of the binding function $\mu(\theta)$. The top three rows show that for all $i \in \{1, 2, 3\}$, the auxiliary parameter $\hat{\mu}_{T,i}$ increases monotonically with the corresponding parameters θ_i of the learning economy, and is much less sensitive to the other parameters θ_j , $j \neq i$ (including σ_δ). Moreover, we note that the auxiliary estimator of b , based on FI ML, is a biased estimator of the parameter b of the incomplete-information economy; this finding illustrates the pitfalls of employing quasi-maximum likelihood estimation in this setting. The bottom row shows that as the noise parameter σ_δ increases, the median return increases monotonically, consistent with the fact that returns become more negatively skewed. In the online appendix, we verify that the third moment is decreasing monotonically with σ_δ . The structural parameter θ is thus well identified by our two candidate auxiliary estimators.

As a benchmark, we also construct a simulated method of moments (SMM) estimator. In the online appendix, we illustrate the impact of the structural parameter θ on the expected values of r_t^n , $n \in \{1, \dots, 4\}$, the leverage coefficient $r_{t-1}r_t^2$, and the volatility autocorrelation measure $r_{t-1}^2r_t^2$. The leverage measure and the second, third and fourth moments appear to be the most sensitive to the structural parameter θ , and are therefore selected for the definition of the SMM estimator.

In Figure 5, we report boxplots of SMM, third moment-based and median-based II estimates of θ obtained from 100 simulated sample paths of length $T = 20,000$ from the learning model with $\bar{k} = 3$ volatility components. For all three estimators, we set the simulation size to $S = 500$, so that each simulated path contains $ST = 10^7$ simulated data points. The indirect inference procedures provide more accurate and less biased estimates of the structural parameters of the learning economy than SMM. The median-based estimator provides substantially more accurate estimates of the parameter σ_δ that controls the agent's information quality. The median-based estimator thus strongly dominates the other two candidate estimators, and we now use it empirically. Overall, the Monte Carlo simulations confirm the excellent properties of the filtering and estimation techniques proposed in the paper.

4.4 Empirical Estimates and Value at Risk Forecasts

We apply our estimation methodology to the daily log excess returns on the U.S. CRSP value-weighted equity index from 2 January 1926 to 31 December 2009. The dataset contains 22,276 observations, which are illustrated in Figure 6. We partition the dataset into an in-sample period, which runs until 31 Dec 1999, and an out-of-sample period, which

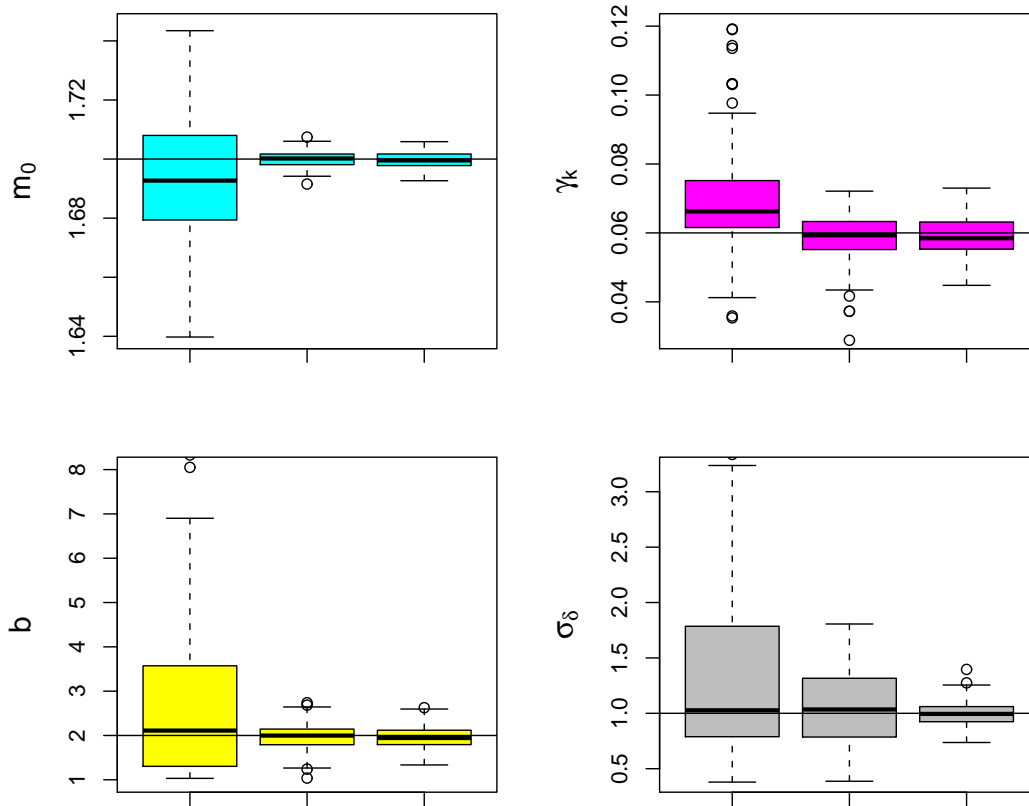


Figure 5: Monte Carlo Simulations of the Learning Model Estimators. This figure illustrates boxplots of the structural parameter estimates obtained using SMM (left boxplot of each panel), the indirect inference estimator based on the third moment (middle boxplots), and the median-based indirect inference estimator (right boxplots). The horizontal lines correspond to the true value of each parameter.




Figure 6: U.S. Equity Return Data. This figure illustrates the daily log excess returns on the CRSP U.S. value-weighted equity index between 2 January 1926 and 31 December 2009. The dashed line separates the in-sample and out-of-sample periods.

covers the remaining ten years.

In Table 3, we report the II estimates of θ . We let $ST = 10^7$ and report standard errors in parentheses. The estimate of σ_δ is significant and declines with \bar{k} .¹⁰ This finding is consistent with the intuition that as \bar{k} increases, the effect of learning becomes increasingly powerful, and a lower σ_δ better matches the negatively skewed excess return series. We also report the log-likelihood of each specification, which is estimated by an SOS filter with $N = 10^7$ particles every period. The likelihood function of the II model increases steadily with \bar{k} . We report in parentheses the t -ratios of a HAC-adjusted Vuong (1989)

¹⁰When $\bar{k} = 1$, the auxiliary parameter is nearly invariant to σ_δ in the relevant region of the parameter space. The Jacobian of the binding function is almost singular, and by (3.7), the estimator of σ_δ has a very large asymptotic variance. The specification with $\bar{k} = 1$ cannot match the median of historical returns and is therefore severely misspecified. These findings illustrate the empirical importance of using higher values of \bar{k} .

Table 3: EMPIRICAL ESTIMATES^a

| \bar{k} | Parameter Estimates | | | | Estimated Likelihood (in logs) |
|-----------|---------------------|--------------------|---------------------|----------------------|--------------------------------|
| | m_0 | $\gamma_{\bar{k}}$ | b | σ_δ | |
| 1 | 1.732 (0.0091) | 0.063 (0.0033) | - | 93.807 (61,616.8) | 65,680.1 (-10.4948) |
| 2 | 1.714 (0.0061) | 0.054 (0.0036) | 21.104 (10.5573) | 4.001 (1.1036) | 67,104.9 (-8.0477) |
| 3 | 1.690 (0.0055) | 0.071 (0.0055) | 16.471 (9.9115) | 2.401 (1.5599) | 67,534.7 (-8.6697) |
| 4 | 1.587 (0.0059) | 0.047 (0.0049) | 5.089 (0.5387) | 1.411 (0.1714) | 68,167.8 |

^aWe report empirical estimates of the learning model (with standard errors in parentheses) based on the daily excess returns of the CRSP index between 2 January 1926 and 31 December 1999. The log-likelihood estimates are based on an SOS filter containing $N = 10^7$ particles. HAC-adjusted Vuong tests comparing $\bar{k} \leq 3$ specifications to $\bar{k} = 4$ are reported in parentheses below the log-likelihood estimates.

test, that is the rescaled differences between the log-likelihoods of the lower-dimensional ($\bar{k} \in \{1, 2, 3\}$) and the highest-dimensional ($\bar{k} = 4$) specifications. The four-component model has a significantly higher likelihood than the other specifications and is therefore selected for the out-of-sample analysis.

We now turn to the out-of-sample implications of the incomplete-information model. The value at risk VaR_{t+1}^p constructed on day t is such that the return on day $t + 1$ will be lower than $-VaR_{t+1}^p$ with probability p . The failure rate is specified as the fraction of observations where the actual return exceeds the value at risk. In a well specified VaR model, the failure rate is on average equal to p . We use as a benchmark historical simulations (e.g. Christoffersen 2009) and Student GARCH(1,1), which are widely used in practice. The historical VaR estimates are based on a window of 60 days, which corresponds to a calendar period of about three months. In Table 4, we report the failure rates of the VaR_{t+1}^p forecasts for $p = 1\%$, 5% , 10% , at horizons of 1 and 5 days produced by: historical simulations, GARCH, the full-information model and the learning model with $\bar{k} = 4$. Standard deviations are reported in parentheses. A failure rate is in bold characters if it differs from its theoretical value at the 1% significance level.

Table 4: FAILURE RATES OF VALUE-AT-RISK FORECASTS^a

| Models | One Day | | | Five Days | | |
|-------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | 1% | 5% | 10% | 1% | 5% | 10% |
| Historical VaR | — | 0.069 (0.0051) | 0.119 (0.0065) | — | 0.066 (0.0111) | 0.129 (0.0150) |
| GARCH | 0.081 (0.0054) | 0.154 (0.0072) | 0.197 (0.0079) | 0.048 (0.0095) | 0.123 (0.0147) | 0.165 (0.0166) |
| $FI, \bar{k} = 4$ | 0.016 (0.0025) | 0.070 (0.0051) | 0.132 (0.0067) | 0.012 (0.0048) | 0.068 (0.0112) | 0.143 (0.0156) |
| $II, \bar{k} = 4$ | 0.008 (0.0018) | 0.047 (0.0042) | 0.094 (0.0058) | 0.014 (0.0052) | 0.060 (0.0106) | 0.135 (0.0153) |

^aThis table reports the failure rates of the 1-day and 5-day value at risk forecasts produced by various methods in the out-of-sample period (2000-2009). The historical VaR is based on a rolling window of 60 days. The GARCH, FI and II forecasts are computed using in-sample parameter estimates. II forecasts are based on an SOS filter with $N = 10^7$ elements. The significance level is 1%.

Historical simulations provide inaccurate VaR forecasts at the 1-day horizon. The failure rates are significantly higher than their theoretical values, which suggests that historical simulations provide overly optimistic estimates of value at risk. GARCH VaR estimates are significantly higher in all cases, while the FI model’s VaR predictions are rejected in three out of six cases. On the other hand, the VaR predictions from the learning model are all consistent with the data. Our empirical findings suggest that the learning model captures well the dynamics of daily stock returns, and outperforms out of sample some of the best reduced-form specifications. We note that this is an excellent result for a consumption-based asset pricing model.

5 Conclusion

In this paper, we have developed powerful filtering and estimation methods for a wide class of learning environments. The new SOS algorithm applies to general state space models in which state-observation pairs can be conveniently simulated. Our method makes no assumption on the availability of the observation density and therefore expands the scope of sequential Monte Carlo methods. The rate of convergence does not depend on the size of the state space, which shows that our filter defeats a form of the curse of

dimensionality. Among many possible applications, SOS is useful to estimate the likelihood function, conduct likelihood-based specification tests, and generate forecasts.

The new filter naturally applies to nonlinear economies with agent learning of the type often considered in financial economics. In this context, SOS permits to track in real time both fundamentals and agent beliefs about fundamentals. Estimation can proceed by simulated ML, but this approach can be computationally costly, as in the example of section 4. For this reason, we have defined an indirect inference estimator by expanding the full-information MLE with a set of statistics that agent learning is designed to capture.

These methods have been applied to a consumption-based asset pricing model with investor learning about multifrequency volatility. We have verified by Monte Carlo simulations the accuracy of our SOS filter and indirect inference estimators, and have implemented these techniques on a long series of daily excess stock returns. We have estimated the parameters driving fundamentals and the quality of the signals received by investors, tracked fundamentals and investor beliefs over time, and verified that the inferred specification provides good value-at-risk forecasts out of sample.

The paper opens multiple directions for future research. SOS can be used to price complex instruments, such as derivatives contracts, which crucially depend on the distribution of the hidden state. We can expand the role of learning in the analysis, for instance by letting the agent learn the parameter of the economy over time, or by conducting the joint online estimation of the structural parameter θ and the state of the economy s_t , as in Polson, Stroud, and Mueller (2008) and Storvik (2002). Further extensions could include inference for equilibrium models with asymmetric information (e.g. Biais, Bossaerts, and Spatt, 2010), and the development of value-at-risk models that incorporate the cross-sectional dispersion of investor beliefs.

A Convergence of the SOS Filter (Section 2)

A.1 A Preliminary Result

In this appendix, we show the convergence of the SOS particle filter defined in section 2 as the number of particles N goes to infinity. Since the path R_T is fixed, our focus is on simulation noise, and expectations in this section are over all the realizations of the random particle method. We begin by establishing the following result for a given $N \geq 1$ and $t \geq 1$.

Lemma A1. *Assume that there exists $U_{t-1}(N)$ such that for every bounded measurable function $\Phi : \mathcal{S} \rightarrow \mathbb{R}$,*

$$\mathbb{E} \left\{ \left[\frac{1}{N} \sum_{n=1}^N \Phi(s_{t-1}^{(n)}) - \mathbb{E}[\Phi(s_{t-1}) | R_{t-1}] \right]^2 \right\} \leq U_{t-1}(N) \|\Phi\|^2. \quad (\text{A.1})$$

Let $U_t^*(N) = 2\kappa_t'^2 A(K)^2 h_t^4 + B(K)\kappa_t / (Nh_t^{nR}) + 2U_{t-1}(N)\kappa_t^2$. Then, the inequality

$$\mathbb{E} \left\{ \left[\frac{1}{N} \sum_{n=1}^N \Phi(\tilde{s}_t^{(n)}) K_{h_t}(r_t - \tilde{r}_t^{(n)}) - f_R(r_t | R_{t-1}) \mathbb{E}[\Phi(s_t) | R_t] \right]^2 \right\} \leq U_t^*(N) \|\Phi\|^2$$

holds for every bounded measurable function Φ .

Proof of Lemma A1. We consider the function

$$a_{t-1}(s_{t-1}) = \int \Phi(\tilde{s}_t) K_{h_t}(r_t - \tilde{r}_t) g(d\tilde{s}_t, d\tilde{r}_t | s_{t-1}, R_{t-1}).$$

We note that

$$\begin{aligned} |a_{t-1}(s_{t-1})| &\leq \|\Phi\| \int K_{h_t}(r_t - \tilde{r}_t) g(d\tilde{s}_t, d\tilde{r}_t | s_{t-1}, R_{t-1}) \\ &= \|\Phi\| \int K_{h_t}(r_t - \tilde{r}_t) f_R(\tilde{r}_t | s_{t-1}, R_{t-1}) d\tilde{r}_t. \end{aligned}$$

The function a_{t-1} is therefore bounded above by $\kappa_t \|\Phi\|$.

The difference $Z = N^{-1} \sum_{n=1}^N \Phi(\tilde{s}_t^{(n)}) K_{h_t}(r_t - \tilde{r}_t^{(n)}) - f_R(r_t | R_{t-1}) \mathbb{E}[\Phi(s_t) | R_t]$ is the

sum of the following three terms:

$$\begin{aligned} Z_1 &= \frac{1}{N} \sum_{n=1}^N \left[\Phi(\tilde{s}_t^{(n)}) K_{h_t}(r_t - \tilde{r}_t^{(n)}) - a_{t-1}(s_{t-1}^{(n)}) \right], \\ Z_2 &= \frac{1}{N} \sum_{n=1}^N a_{t-1}(s_{t-1}^{(n)}) - \int a_{t-1}(s_{t-1}) \lambda(ds_{t-1} | R_{t-1}), \\ Z_3 &= \int a_{t-1}(s_{t-1}) \lambda(ds_{t-1} | R_{t-1}) - f_R(r_t | R_{t-1}) \mathbb{E}[\Phi(s_t) | R_t]. \end{aligned}$$

Let $\mathcal{S}_{t-1}^{(N)} = (s_{t-1}^{(1)}, \dots, s_{t-1}^{(N)})$ denote the vector of period $-(t-1)$ particles. Conditional on $\mathcal{S}_{t-1}^{(N)}$, Z_1 has a zero mean, while Z_2 and Z_3 are deterministic. Hence:

$$\mathbb{E}(Z^2) = \mathbb{E}(Z_1^2) + \mathbb{E}[(Z_2 + Z_3)^2] \leq \mathbb{E}(Z_1^2) + 2\mathbb{E}(Z_2^2) + 2\mathbb{E}(Z_3^2).$$

Conditional on $\mathcal{S}_{t-1}^{(N)}$, the state-observation pairs $\{(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})\}_{n=1}^N$ are independent, and each $(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})$ is drawn from $g(\cdot | s_{t-1}^{(n)}, R_{t-1})$; the addends of $\Phi(\tilde{s}_t^{(n)}) K_{h_t}(r_t - \tilde{r}_t^{(n)}) - a_{t-1}(s_{t-1}^{(n)})$ are thus independent and have mean zero. We infer that the conditional expectation of Z_1^2 is bounded above by:

$$\frac{1}{N^2} \sum_{n=1}^N \int \Phi(\tilde{s}_t)^2 K_{h_t}(r_t - \tilde{r}_t)^2 g(d\tilde{s}_t, d\tilde{r}_t | s_{t-1}^{(n)}, R_{t-1}) \leq \frac{\kappa_t \|\Phi\|^2}{N} \int K_{h_t}(r_t - \tilde{r}_t)^2 d\tilde{r}_t.$$

We apply the change of variable $u = (r_t - \tilde{r}_t)/h_t$:

$$\int K_{h_t}(r_t - \tilde{r}_t)^2 d\tilde{r}_t = \frac{B(K)}{h_t^{n_R}},$$

and infer that $\mathbb{E}(Z_1^2) \leq \|\Phi\|^2 B(K) \kappa_t / (N h_t^{n_R})$.

Since the function $a_{t-1}(s_{t-1})$ is bounded above by $\kappa_t \|\Phi\|$, we infer from (A.1) that: $\mathbb{E}(Z_2^2) \leq U_{t-1}(N) \kappa_t^2 \|\Phi\|^2$.

Finally, we observe that $f_R(r_t | R_{t-1}) \mathbb{E}[\Phi(s_t) | R_t] = \int \Phi(s_t) f_R(r_t | s_t, R_{t-1}) \lambda(ds_t | R_{t-1})$, and therefore

$$\begin{aligned} Z_3 &= \int \Phi(s_t) \left\{ \int K_{h_t}(r_t - \tilde{r}_t) [f_R(\tilde{r}_t | s_t, R_{t-1}) - f_R(r_t | s_t, R_{t-1})] d\tilde{r}_t \right\} \lambda(ds_t | R_{t-1}) \\ &= \int \Phi(s_t) \left\{ \int K(u) [f_R(r_t - h_t u | s_t, R_{t-1}) - f_R(r_t | s_t, R_{t-1})] du \right\} \lambda(ds_t | R_{t-1}). \end{aligned}$$

Note that $|\int K(u)[f_R(r_t - h_t u|s_t, R_{t-1}) - f_R(r_t|s_t, R_{t-1})]du| \leq \kappa'_t A(K)h_t^2$. Hence $|Z_3| \leq \kappa'_t A(K)h_t^2 \|\Phi\|$ and therefore $\mathbb{E}(Z_3^2) \leq \kappa_t'^2 A(K)^2 h_t^4 \|\Phi\|^2$. We conclude that the lemma holds. *Q.E.D.*

A.2 Proof of Theorem 4

The proof of (2.5) proceeds by induction. When $t = 0$, the particles are drawn from the prior λ_0 , and the conditional expectation is computed under the same prior. Hence the property (2.5) holds with $U_0(N) = 1/N$.

We now assume that the property (2.5) holds at date $t - 1$. The estimation error $X = N^{-1} \sum_{n=1}^N \Phi(s_t^{(n)}) - \mathbb{E}[\Phi(s_t)|R_t]$ is the sum of:

$$\begin{aligned} X_1 &= \frac{1}{N} \sum_{n=1}^N \Phi(s_t^{(n)}) - \sum_{n=1}^N p_t^{(n)} \Phi(\tilde{s}_t^{(n)}). \\ X_2 &= \left[\sum_{n=1}^N p_t^{(n)} \Phi(\tilde{s}_t^{(n)}) \right] \left[\frac{f_R(r_t|R_{t-1}) - N^{-1} \sum_{n'=1}^N K_{h_t}(r_t - \tilde{r}_t^{(n')})}{f_R(r_t|R_{t-1})} \right], \\ X_3 &= \frac{1}{N f_R(r_t|R_{t-1})} \sum_{n=1}^N \Phi(\tilde{s}_t^{(n)}) K_{h_t}(r_t - \tilde{r}_t^{(n)}) - \mathbb{E}[\Phi(s_t)|R_t]. \end{aligned}$$

The first term, X_1 , corresponds to step 3 resampling, the second term to the normalization of the resampling weights, and the third term to the error in the estimation of Φ using the nonnormalized weights.

Conditional on $\{(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})\}_{n=1}^N$, the particles $s_t^{(n)}$ are independent and identically distributed, and X_1 has mean zero. We infer that $\mathbb{E}[X_1^2 | \{\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)}\}_{n=1}^N] \leq \|\Phi\|^2/N$, and therefore $\mathbb{E}(X_1^2) \leq \|\Phi\|^2/N$. Note that when we use stratified, residual or combined stratified-residual resampling in step 3, the inequality $\mathbb{E}(X_1^2) \leq \|\Phi\|^2/N$ remains valid, and smaller upper bounds can also be derived.¹¹

Conditional on $\{(\tilde{s}_t^{(n)}, \tilde{r}_t^{(n)})\}_{n=1}^N$, X_2 and X_3 are deterministic variables. The mean squared error satisfies:

$$\mathbb{E}(X^2) = \mathbb{E}(X_1^2) + \mathbb{E}[(X_2 + X_3)^2] \leq \mathbb{E}(X_1^2) + 2\mathbb{E}(X_2^2) + 2\mathbb{E}(X_3^2).$$

We note that $|X_2| \leq \|\Phi\| [f_R(r_t|R_{t-1})]^{-1} \left| f_R(r_t|R_{t-1}) - \sum_{n'=1}^N K_{h_t}(r_t - \tilde{r}_t^{(n')})/N \right|$. Us-

¹¹See Cappé, O., Moulines, E., and T. Rydén (2005, ch. 7) for a detailed discussion of sampling variance.

ing the induction hypothesis at date $t - 1$, we apply Lemma A1 with $\Phi \equiv 1$ and obtain that $\mathbb{E}(X_2^2)$ is bounded above by:

$$\frac{U_t^*(N)\|\Phi\|^2}{[f_R(r_t|R_{t-1})]^2}. \quad (\text{A.2})$$

Lemma A1 implies that $\mathbb{E}(X_3^2)$ is also bounded above by (A.2). We conclude that $\mathbb{E}(X^2) \leq U_t(N)\|\Phi\|^2$, where $U_t(N) = 4U_t^*(N)[f_R(r_t|R_{t-1})]^{-2} + N^{-1}$, or equivalently

$$U_t(N) = \frac{4}{[f_R(r_t|R_{t-1})]^2} \left[2\kappa_t'^2 A(K)^2 h_t^4 + \frac{B(K)\kappa_t}{N h_t^{n_R}} + 2U_{t-1}(N)\kappa_t^2 \right] + \frac{1}{N}. \quad (\text{A.3})$$

This establishes part (2.5) of the theorem. From (2.5) and Lemma A1 with $\Phi \equiv 1$, (2.4) follows.

Assume now that the bandwidth is a function of N , and that assumption 3 holds. A simple recursion implies that $\lim_{N \rightarrow \infty} U_t(N) = 0$ for all t . The mean squared error converges to zero for any bounded measurable function Φ .

We now characterize the rate of convergence. Given $U_{t-1}(N)$, we know that the coefficient $U_t(N)$ defined by (A.3) is minimal if

$$h_t = N^{-1/(n_R+4)} \left[\frac{\kappa_t n_R B(K)}{8\kappa_t'^2 A(K)^2} \right]^{1/(n_R+4)}. \quad (\text{A.4})$$

More generally, if the bandwidth sequence is of the form $h_t(N) = h_t(1)/N^{-1/(n_R+4)}$, then $U_t(N)$ is of the form:

$$U_t(N) = u_{1,t}N^{-4/(n_R+4)} + u_{2,t}U_{t-1}(N) + N^{-1}. \quad (\text{A.5})$$

where $u_{1,t}$ and $u_{2,t}$ are finite nonnegative coefficients.¹² By a simple recursion, $U_t(N)$ is of order $N^{-4/(n_R+4)}$ for all t . *Q.E.D.*

¹²We verify that $u_{1,t} = 4[f(r_t|R_{t-1})]^{-2} [2\kappa_t'^2 h_t(1)^4 A(K)^2 + B(K)\kappa_t h_t(1)^{-n_R}]$ and $u_{2,t} = 8\kappa_t^2 [f(r_t|R_{t-1})]^{-2}$.

B Learning Economies (Section 3)

B.1 Proof of Proposition 6

We infer from Bayes' rule that

$$\Pi_t^j \propto \underbrace{f_X(x_t | M_t = m^j, X_{t-1}; \theta)}_{=f_X(x_t | M_t = m^j; \theta) \text{ by As. 5(b)}} \mathbb{P}(M_t = m^j | X_{t-1}; \theta),$$

where

$$\mathbb{P}(M_t = m^j | X_{t-1}; \theta) = \sum_{i=1}^d \underbrace{\mathbb{P}(M_t = m^j | M_{t-1} = m^i, X_{t-1}; \theta)}_{=a_{ij} \text{ by As. 5(a)}} \mathbb{P}(M_{t-1} = m^i | X_{t-1}; \theta),$$

and Proposition 6 holds. *Q.E.D.*

B.2 Proof of Proposition 7

Bayes' rule (3.1) implies that for every $i \in \{1, \dots, d\}$,

$$\Pi_t | M_t = m^i, s_{t-1}, \dots, s_1 \sim \Pi_t | M_t = m^i, \Pi_{t-1}. \quad (\text{B.1})$$

Also, by Assumption 5(a)

$$\mathbb{P}(M_t = m^i | s_{t-1}, \dots, s_1; \theta) = \mathbb{P}(M_t = m^i | M_{t-1}; \theta). \quad (\text{B.2})$$

From (B.1) and (B.2), we conclude that s_t is first-order Markov.

We know from Kaijser (1975) that under the conditions stated in the proposition, the belief process Π_t has a unique invariant distribution. Proposition 2.1 in van Handel (2009) implies that (M_t, Π_t) also has a unique invariant measure Λ_∞ .¹³ We infer from the Birkhoff-Khinchin theorem that for any integrable function $\Phi : \mathcal{S} \rightarrow \mathbb{R}$, the sample average $T^{-1} \sum_{t=1}^T \Phi(s_t)$ converges almost surely to the expectation of Φ under the invariant measure Λ_∞ . *Q.E.D.*

¹³Chigansky (2006) derives a similar result in continuous time.

B.3 Proof of Proposition 9

The econometrician recursively applies Bayes' rule:

$$\mathbb{P}(M_t = m^j | R_t; \phi) = \frac{f_{R,FI}(r_t | M_t = m^j, R_{t-1}; \phi) \mathbb{P}(M_t = m^j | R_{t-1}; \phi)}{f_{R,FI}(r_t | R_{t-1}; \phi)}$$

Since $f_{R,FI}(r_t | M_t = m^j, R_{t-1}; \phi) = \sum_{i=1}^d f_{i,j}(r_t; \phi) \mathbb{P}(M_{t-1} = m^i | M_t = m^j, R_{t-1}; \phi)$, we infer that $f_{R,FI}(r_t | M_t = m^j, R_{t-1}; \phi) \mathbb{P}(M_t = m^j | R_{t-1}; \phi) = \sum_{i=1}^d f_{i,j}(r_t; \phi) \mathbb{P}(M_{t-1} = m^i, M_t = m^j | R_{t-1}; \phi)$, and therefore

$$\mathbb{P}(M_t = m^j | R_t; \phi) = \frac{\sum_{i=1}^d a_{i,j} f_{i,j}(r_t; \phi) \mathbb{P}(M_{t-1} = m^i | R_{t-1}; \phi)}{f_{R,FI}(r_t | R_{t-1}; \phi)}.$$

The econometrician's conditional probabilities are therefore computed recursively.

Since the conditional probabilities $\mathbb{P}(M_t = m^j | R_t; \phi)$ add up to unity, the conditional density of r_t satisfies

$$f_{R,FI}(r_t | R_{t-1}; \phi) = \sum_{i=1}^d \sum_{j=1}^d a_{i,j} f_{i,j}(r_t; \phi) \mathbb{P}(M_{t-1} = m^i | R_{t-1}; \phi).$$

The log-likelihood function $\mathcal{L}_{FI}(\phi | R_T) = \sum_{t=1}^T \ln f_{R,FI}(r_t | R_{t-1}; \phi)$ thus has an analytical expression. *Q.E.D.*

B.4 Indirect Inference Estimator

We provide a set of sufficient conditions for the asymptotic results at the end of section 3.2, and then discuss numerical implementation.

B.4.1 Sufficient Conditions for Convergence

We assume that $\hat{\eta}_T$ maximizes a criterion $\mathcal{H}(\eta, R_T)$ that does not depend on the full-information MLE $\hat{\phi}_T$. The auxiliary estimator $\hat{\mu}_T = (\hat{\phi}_T', \hat{\eta}_T')'$ can therefore be written as:

$$\hat{\mu}_T = \arg \max_{\mu} \mathcal{Q}_T(\mu, R_T), \tag{B.3}$$

where $\mathcal{Q}_T(\mu, R_T) = T^{-1} \mathcal{L}_{FI}(\phi, R_T) + \mathcal{H}(\eta, R_T)$ for all $\mu = (\phi', \eta)'$.

Assumption 10 (Binding Function) Under the structural model θ^* , the auxiliary criterion function $\mathcal{Q}_T(\mu, R_T)$ converges in probability to $\mathcal{Q}_\infty(\mu, \theta^*)$ for all μ . Moreover, the function $\mu : \mathbb{R}^p \rightarrow \mathbb{R}^p$ defined by

$$\mu(\theta) = \arg \max_{\mu} \mathcal{Q}_\infty(\mu, \theta),$$

called the binding function, is injective.

Assumption 11 (Score) The renormalized score satisfies:

$$\sqrt{T} \frac{\partial \mathcal{Q}_T}{\partial \mu} [\mu(\theta^*), R_T] \xrightarrow{d} \mathcal{N}(0, I_0),$$

where I_0 is positive definite symmetric matrix.

Assumption 12 (Hessian of Criterion Function) The Hessian matrix

$$-\frac{\partial^2 \mathcal{Q}_T}{\partial \mu \partial \mu'} [\mu(\theta^*), R_T]$$

is invertible and converges in probability to a nonsingular matrix J_0 .

Under assumptions 10-12, the auxiliary estimator satisfies $\sqrt{T} [\hat{\mu}_T - \mu(\theta^*)] \xrightarrow{d} \mathcal{N}(0, W^*)$, where $W^* = J_0^{-1} I_0 J_0^{-1}$, and the asymptotic results at the end of section 3.2 hold (Gouriéroux, Monfort, and Renault, 1993; Gouriéroux and Monfort, 1996).

B.4.2 Numerical Implementation

Since in the just-identified case $\hat{\mu}_{ST}(\hat{\theta}_T) = \hat{\mu}_T$, the simulated auxiliary estimator $\hat{\mu}_{ST}(\theta)$ satisfies

$$\frac{\partial Q_{ST}}{\partial \mu} [\underbrace{\hat{\mu}_{ST}(\theta)}_{\hat{\mu}_T}, R_{ST}(\theta)] = 0.$$

Hence, the indirect inference estimator $\hat{\theta}_T$ minimizes the EMM-type objective function:

$$\left\{ \frac{\partial Q_{ST}}{\partial \mu} [\hat{\mu}_T, R_{ST}(\theta)] \right\}' W_T \left\{ \frac{\partial Q_{ST}}{\partial \mu} [\hat{\mu}_T, R_{ST}(\theta)] \right\}, \quad (\text{B.4})$$

where W_T is any positive-definite weighting matrix. This property can be used to compute $\hat{\theta}_T$. For each iteration of θ , the evaluation of the EMM objective function (B.4) requires only the evaluation of the score. By contrast, the evaluation of the objective function (3.6) requires the optimization of the FI likelihood in order to obtain $\hat{\mu}_{ST}(\theta)$. The computational advantage of EMM is substantial in applications where the calculation of the full-information MLE is expensive.

In the just-identified case and under assumptions 10-13, the asymptotic variance-covariance matrix of the indirect inference estimator simplifies to

$$\Sigma = \left(1 + \frac{1}{S}\right) \left\{ \frac{\partial^2 \mathcal{Q}_\infty}{\partial \theta \partial \mu'} [\mu(\theta^*), \theta^*] I_0^{-1} \frac{\partial^2 \mathcal{Q}_\infty}{\partial \mu \partial \theta'} [\mu(\theta^*), \theta^*] \right\}^{-1},$$

as shown in Gouriéroux and Monfort (1996). Note that the choice of the weighting matrix W_T does not affect the asymptotic variance of the indirect inference estimator in the exactly identified case.

In practice, we can estimate I_0 and $\frac{\partial^2 \mathcal{Q}_\infty}{\partial \theta \partial \mu'} [\mu(\theta^*), \theta^*]$ in the following way.

Assumption 13 (Decomposable Score) *The score function can be written as:*

$$\frac{\partial \mathcal{Q}_T}{\partial \mu}(\mu, R_T) \equiv \frac{1}{T} \sum_{t=1}^T \psi(r_t | R_{t-1}; \mu)$$

for all R_T and μ .

Note that Assumption 13 is satisfied by the median-based and the third moment-based indirect inference estimators considered in section 4.

By Assumption 13, the auxiliary parameter satisfies the first-order condition:

$$\frac{\partial \mathcal{Q}_T}{\partial \mu}(\hat{\mu}_T, R_T) = \frac{1}{T} \sum_{t=1}^T \psi(r_t | R_{t-1}; \hat{\mu}_T) = 0. \quad (\text{B.5})$$

We estimate I_0 by the Newey and West (1987) variance-covariance matrix:

$$\hat{I}_0 = \hat{\Gamma}_0 + \sum_{v=1}^{\tau} \left(1 - \frac{v}{\tau + 1}\right) \left(\hat{\Gamma}_v + \hat{\Gamma}_v'\right), \quad (\text{B.6})$$

where $\hat{\Gamma}_v = T^{-1} \sum_{t=v+1}^T \psi(r_t | R_{t-1}; \hat{\mu}_T) \psi(r_t | R_{t-1}; \hat{\mu}_T)'$. All the results reported in the paper

are based on $\tau = 10$ lags. We approximate $\frac{\partial^2 Q_\infty}{\partial\theta\partial\mu'}[\mu(\theta^*), \theta^*]$ by

$$\frac{\partial^2 Q_{ST}}{\partial\theta\partial\mu'}[\hat{\mu}_T, R_{ST}(\hat{\theta}_T)],$$

and obtain a finite-sample estimate of the asymptotic variance-covariance matrix Σ .

References

- ANDERSEN, T.G., H.J. CHUNG AND B.E. SORENSEN (1999). Efficient Method of Moment Estimation of a Stochastic Volatility Model: A Monte Carlo Study. *Journal of Econometrics* **91**, 61–87.
- ANDRIEU, C., AND A. DOUCET (2002). Particle Filtering for Partially Observed Gaussian State Space Models. *Journal of the Royal Statistical Society B* **64(4)**, 827–836.
- ANDRIEU, C., DOUCET, A., AND R. HOLENSTEIN (2010). Particle Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society B* **72(3)**, 269–342.
- BIAIS, B., BOSSAERTS, P., AND C. SPATT (2010). Equilibrium Asset Pricing and Portfolio Choice under Asymmetric Information. Forthcoming, *Review of Financial Studies*.
- BRANDT, M. W., ZENG, Q., AND L. ZHANG (2004). Equilibrium Stock Return Dynamics under Alternative Rules of Learning about Hidden States. *Journal of Economic Dynamics and Control* **28**, 1925–1954.
- BRENNAN, M. (1998). The Role of Learning in Dynamic Portfolio Decisions. *European Finance Review* **1**, 295–396.
- BRENNAN, M., AND Y. XIA (2001). Stock Price Volatility and the Equity Premium. *Journal of Monetary Economics* **47**, 249–283.
- CALVET, L. E., AND A. J. FISHER (2001). Forecasting Multifractal Volatility. *Journal of Econometrics* **105**, 27–58.
- CALVET, L. E., AND A. J. FISHER (2004). How to Forecast Long-Run Volatility: Regime-Switching and the Estimation of Multifractal Processes. *Journal of Financial Econometrics* **2**, 49–83.
- CALVET, L. E., AND A. J. FISHER (2007). Multifrequency News and Stock Returns. *Journal of Financial Economics* **86**: 178-212.
- CALVET, L. E., AND A. J. FISHER (2008). *Multifractal Volatility: Theory, Forecasting and Pricing*. Elsevier – Academic Press.
- CALVET, L. E., FISHER, A. J., AND S. THOMPSON (2006). Volatility Comovement: A Multifrequency Approach. *Journal of Econometrics* **131**, 179–215.

- CALZOLARI G., FIORENTINI, G. AND SENTANA, E. (2004). Constrained Indirect Estimation. *Review of Economic Studies* **71**, 945–973.
- CAPPÉ, O., MOULINES, E., AND T. RYDÉN (2005). *Inference in Hidden Markov Models*. Springer.
- CECCHETTI, S. G., LAM, P. AND MARK, N. C. (2000). Asset Pricing with Distorted Beliefs: Are Equity Returns Too Good to Be True? *American Economic Review* **90(4)**, 787–805.
- CHIB, S., NARDARI, F., AND N. SHEPHARD (2006). Analysis of High-Dimensional Multivariate Stochastic Volatility Models. *Journal of Econometrics* **134**, 341–371.
- CHIGANSKY, P. (2006). An Ergodic Theorem for Filtering with Applications to Stability. *Systems and Control Letters* **55(11)**, 908–917.
- CHOPIN, N. (2004). Central Limit Theorem for Sequential Monte Carlo Methods and its Applications to Bayesian Inference. *Annals of Statistics* **32(6)**, 2385–2411.
- CHRISTOFFERSEN, P. (2009). Backtesting. In: *Encyclopedia of Quantitative Finance*, R. Cont (ed)., John Wiley and Sons.
- CHRISTOFFERSEN, P., JACOBS, K., AND K. MIMOUNI (2007). Volatility Dynamics for the S&P 500: Evidence from Realized Volatility, Daily Returns and Option Prices. Forthcoming, *Review of Financial Studies*.
- CRISAN, D., AND DOUCET, A. (2002). A Survey of Convergence Results on Particle Filtering Methods for Practitioners. *IEEE Transactions on Signal Processing* **50(3)** 736–746.
- CZELLAR, V., KAROLYI, G. A. AND RONCHETTI, E. (2007). Indirect Robust Estimation of the Short-Term Interest Process. *Journal of Empirical Finance* **14**, 546–563.
- CZELLAR, V. AND RONCHETTI E. (2010). Accurate and Robust Tests for Indirect Inference. *Biometrika* **97(3)**, 621–630.
- DAVID, A. (1997). Fluctuating Confidence in Stock Markets: Implications for Returns and Volatility. *Journal of Financial and Quantitative Analysis* **32**, 427–462.

- DAVID, A., AND P. VERONESI (2006). Inflation and Earnings Uncertainty and Volatility Forecasts. Working paper, University of Calgary and University of Chicago.
- DEL MORAL, P. (2004). *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer-Verlag, New York.
- DOUCET, A., AND A. M. JOHANSEN (2008). A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later. Working paper, The Institute of Statistical Mathematics and the University of Warwick.
- DRIDI, R., GUAY, A. AND RENAULT, E. (2007). Indirect Inference and Calibration of Dynamic Stochastic General Equilibrium Model. *Journal of Econometrics* **136**, 397–430.
- FEARNHEAD, P., AND P. CLIFFORD (2003). On-line Inference for Hidden Markov Models via Particle Filters. *Journal of the Royal Statistical Society B* **65**, 887899.
- FERNÁNDEZ-VILLAVERDE, J., GUERRÓN-QUINTANA, P., RUBIO-RAMIREZ, J., AND M. URIBE (2009). Risk Matters: The Real Effects of Volatility Shocks. Forthcoming in *American Economic Review*.
- FERNÁNDEZ-VILLAVERDE, J., AND J. RUBIO-RAMIREZ (2007). Estimating Macroeconomic Models: A Likelihood Approach. *Review of Economic Studies* **74**, 1059–1087.
- GENTON, M. G., AND E. RONCHETTI (2003). Robust Indirect Inference. *Journal of the American Statistical Association* **98**, 67–76.
- GILKS, W. R., AND C. BERZUINI (2001). Following a Moving Target – Monte Carlo Inference for Dynamic Bayesian Models. *Journal of the Royal Statistical Society B* **63**, 127–146.
- GODSILL, S. J., DOUCET, A., AND M. WEST (2004). Monte Carlo Smoothing for Non-linear Time Series. *Journal of the American Statistical Association* **99**, 156–168.
- GORDON, N., SALMOND, D., AND A. F. SMITH (1993). Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation. *IEE Proceedings F* **140**, 107–113.
- GOURIÉROUX, C., MONFORT, A. AND RENAULT, E. (1993). Indirect Inference. *Journal of Applied Econometrics* **8**, S85–S118.

- GOURIÉROUX, C. AND MONFORT, A. (1996). *Simulation-Based Econometric Methods*. – Oxford University Press.
- GUIDOLIN, M. AND TIMMERMANN, A. (2003). Option Prices under Bayesian Learning: Implied Volatility Dynamics and Predictive Densities. *Journal of Economic Dynamics and Control* **27**, 717–769.
- HAMILTON, J. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* **57**, 357–384.
- HANSEN, L. P. (2007). Beliefs, Doubts and Learning: Valuing Macroeconomic Risk. *American Economic Review* **97**(2), 1–30.
- HANSEN, L., POLSON, N., AND T. SARGENT (2011). Nonlinear Filtering and Learning Dynamics. Forthcoming, *Journal of Business and Economic Statistics*.
- HEGGLAND, K. AND FRIGESSI, A. (2004). Estimating Functions in Indirect Inference. *Journal of the Royal Statistical Society Series B*, **66**, 447–462.
- JOHANNES, M., AND N. POLSON (2009), “Particle Filtering,” in *Handbook of Financial Time Series*, ed. by T. Andersen, R. A. Davis, J.-P. Kreiss, and Th. Mikosch, 1015–1028. Springer.
- JOHANNES, M., POLSON, N., AND J. STROUD (2009), Optimal Filtering of Jump Diffusions: Extracting Latent States from Asset Prices. *Review of Financial Studies* **22**(7), 2759–2799.
- KAIJSER, T. (1975). A Limit Theorem for Partially Observed Markov Chains. *Annals of Probability* **3**(4), 677–696.
- KIM, S., SHEPHARD, N., AND S. CHIB (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies* **65**(3), 361–393.
- KITAGAWA, G. (1996). Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models. *Journal of Computational and Graphical Statistics* **5**, 1–25.
- KUENSCH, H. R. (2005). Recursive Monte Carlo Filters: Algorithms and Theoretical Analysis. *Annals of Statistics* **33**(5), 1983–2021.

- LETTAU, M., LUDVIGSON, S. AND WACHTER J. (2008). The Declining Equity Premium: What Role Does Macroeconomic Risk Play? *Review of Financial Studies* **21**, 1653–1687.
- LIU, J. L. AND R. CHEN (1998). Sequential Monte Carlo Methods for Dynamic Systems. *Journal of the American Statistical Association* **93**, 1032–1044.
- MOORE, B. AND SCHALLER, H. (1996). Learning, Regime Switches, and Equilibrium Asset Pricing Dynamics. *Journal of Economic Dynamics and Control* **20(6-7)**, 979–1006.
- NEWWEY, W., AND K. WEST (1987). A Simple, Positive Definite Heteroscedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* **55**, 703–708.
- PÁSTOR, L., AND VERONESI, P. (2009a). Learning in Financial Markets. *Annual Review of Financial Economics* **1**, 361 – 381.
- PÁSTOR, L., AND VERONESI, P. (2009b). Technological Revolutions and Stock Prices. *American Economic Review* **99**, 1451–1483.
- PITT, M. (2005). Smooth Particle Filters for Likelihood Evaluation and Maximisation. Working paper, University of Warwick.
- PITT, M., AND N. SHEPHARD (1999). Filtering Via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association* **94**, 590–599.
- POLSON, N., STROUD, J., AND P. MUELLER (2008). Practical Filtering with Sequential Parameter Learning. *Journal of the Royal Statistical Society B* **70(2)**, 413–428.
- ROSSI, V., AND J.-P. VILA (2006). Nonlinear Filtering in Discrete Time: A Particle Convolution Approach. *Annales de l'Institut de Statistique de l'Université de Paris* **50(3)**, 71–102.
- ROSSI, V., AND J.-P. VILA (2009). Convolution Particle Filter for Parameter Estimation in General State-Space Models. *IEEE Transactions on Aerospace and Electronic Systems* **45(3)**, 1063– 1072.
- SMITH, A. A. (1993). Estimating Nonlinear Time Series Models Using Simulated Vector Autoregressions. *Journal of Applied Econometrics* **8**, S63–S84.
- STORVIK, G. (2002). Particle Filters in State Space Models with the Presence of Unknown Static Parameters. *IEEE Transactions on Signal Processing* **50**, 281–289.

- TIMMERMANN, A. G. (1993). How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock prices. *Quarterly Journal of Economics* **108**, 1135–1145.
- TIMMERMANN, A. G. (1996). Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning. *Review of Economic Studies* **63**, 523–557.
- VAN HANDEL, R. (2009). Uniform Time Average Consistency of Monte Carlo Particle Filters. *Stochastic Processes and Their Applications* **119**, 3835–3861.
- VAN NIEUWERBURGH, S. AND VELDKAMP, L. (2006). Learning Asymmetries in Real Business Cycles. *Journal of Monetary Economics* **53**, 753–772.
- VERONESI, P. (1999). Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model. *Review of Financial Studies* **12**, 975–1007.
- VERONESI, P. (2000). How Does Information Quality Affect Stock Returns? *Journal of Finance* **55**, 807–837.
- VUONG, Q. (1989). Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses. *Econometrica* **57**, 307–333.