Econometrics and Statistics 000 (2016) 1–21

Contents lists available at ScienceDirect

Econometrics and Statistics

journal homepage: www.elsevier.com/locate/ecosta

Filterbased stochastic volatility in continuous-time hidden Markov models

Vikram Krishnamurthy[a], Elisabeth Leoff[b], Jörn Sass[c,*]

[a] Department of Electrical and Computer Engineering, Cornell Tech, Cornell University, USA
[b] Department of Financial Mathematics, Fraunhofer Institute for Industrial Mathematics ITWM, Kaiserslautern, Germany
[c] Department of Mathematics, University of Kaiserslautern, Kaiserslautern, Germany

ARTICLE INFO

Article history:
Received 5 February 2016
Revised 14 October 2016
Accepted 14 October 2016
Available online xxx

Keywords:
Markov switching model
Non-constant volatility
Stylized facts
Portfolio optimization
Social learning

ABSTRACT

Regime-switching models, in particular Hidden Markov Models (HMMs) where the switching is driven by an unobservable Markov chain, are widely-used in financial applications, due to their tractability and good econometric properties. In continuous time, properties of HMMs with constant and of HMMs with switching volatility can be quite different. To have a realistic model with unobservable Markov chain in continuous time and good econometric properties, a regime-switching model where the volatility depends on the filter for the underlying chain is introduced and the filtering equations are stated. Such models are motivated by agent based social learning models in economics. An approximation result for a fixed information filtration is proved and further motivation is provided by considering social learning arguments. The relation to the switching volatility model is analyzed in detail and a convergence result for the discretized model is given. Econometric properties are illustrated by numerical simulations.

1. Introduction

Regime-switching models are very popular in the field of mathematical finance to describe return processes with time-changing drift or volatility parameters. They are a possible way to generalize the classical Black–Scholes log normal stock price model by making the parameters dependent on a Markov chain with finitely many states. This Markov chain can be interpreted as encoding all the information in the market into one single process that describes the underlying state of the economy. As such this chain is assumed to be unobservable, the normal investor does not have access to all the possible information. The only observation that can be made is the return process itself.

In this work, we consider both the Hidden Markov Model (HMM) and what we call the Markov Switching Model (MSM) in continuous time. Both models have been studied extensively and are popular for applications. The difference between them is only that the HMM has constant volatility, while in the MSM the volatility jumps with the Markov chain. This difference leads to a difference in the behavior of information between both models, since the volatility is observable via the quadratic variation in continuous time. Such a distinction is not present in discrete time, cf. Elliott et al. (2008), Elliott and Siu (2012), Guo (2001), Siu (2011). For comparison of estimation procedures for the continuous-time MSM we refer to Hahn et al. (2010).

* Corresponding author.
E-mail addresses: vikramk@cornell.edu (V. Krishnamurthy), elisabeth.leoff@itwm.fraunhofer.de (E. Leoff), sass@mathematik.uni-kl.de (J. Sass).

http://dx.doi.org/10.1016/j.ecosta.2016.10.007

In finance, one often needs models that are defined in continuous time since these allow for explicit results. To be applicable these have to be close to the results for the real-world discrete-time setting. For example, in portfolio optimization the aim is to find the strategy that maximizes the expected utility at terminal time. These optimal strategies can be derived explicitly only in continuous time, but in the HMM they are close to the discrete-time strategies, cf. the detailed discussion in Section 2.2. Consequently, the continuous-time HMM is such a model which is very well tractable and yields results close to the discrete-time model. Thus we may call it discretization-consistent.

The volatility in the continuous-time HMM is assumed to be constant. Empirical observations of financial data across a wide range of instruments, markets and time periods indicate that for many real data sets this is not true, as was already noted by e.g., Fama (1965). This leads to the concept of stochastic volatility in model building. There have been many different modeling approaches leading to various models where the volatility itself is a stochastic process. Popular models are e.g., Engle's famous Autoregressive Conditional Heteroskedasticity model (ARCH) from 1982, see Engle (1982), and its generalization GARCH, see Bollerslev (1986) and also cf. Taylor (1994), or in continuous time the well-known Heston model (Heston, 1993).

In the context of regime-switching models, the natural way to introduce a non-constant volatility is to use the Markov chain to control the volatility, as usually done in discrete time when considering regime switching models, see Hamilton (1989). Switching can also be introduced in the volatility models mentioned above. Switching between constant values for the volatility leads in continuous time to the MSM, which has better econometric properties than the HMM for many applications. Some stylized facts such as volatility clustering cannot be reproduced in the HMM. However, it turns out that in the continuous-time MSM the Markov chain can be observed using the quadratic variation of the return process, see Proposition 2.1 below. Therefore, results obtained in the MSM can be far away from results obtained in the corresponding discrete-time model. Thus we do not have the discretization-consistency for the continuous-time MSM as we have it for the HMM.

To get a model which is both consistent and has good econometric properties, the idea now is to introduce stochastic volatility in a continuous-time HMM. Haussmann and Sass (2004) propose a regime-switching model where the volatility depends on an observable diffusion. They prove that despite the non-constant volatility in this model the chain is still hidden and derive its filtering equation which is again finite-dimensional. The question of how to actually model the volatility process is still left open. Due to the better econometric properties of the MSM our goal is to find a volatility process that leads to an HMM that in a sense approximates the continuous-time MSM.

Thus, in this work, we introduce an HMM in continuous time where the volatility is a linear function of the (normalized) filter. For a given filtration we construct the volatility process adapted to this filtration that results in the best approximation of the MSM-returns in the mean squared sense. This result motivates the introduction of the Filterbased-Volatility Hidden Markov Model (FB-HMM), where the volatility is a linear function of the normalized filter. In this model, the returns and the filter constitute a system of equations coupled by the volatility process. In a sense, this model lies between the HMM and MSM: the chain is still hidden as in the HMM, but the volatility is non-constant similar to the MSM. The function that connects the filter to the volatility process in the FB-HMM is even the same function that connects the chain to the volatility in the MSM. Another motivation for the model comes from an instantaneous iteration of observing the returns and trading (resulting in a new volatility) in a financial market, which is similar to a social learning model used in economics, see Remark 3.3.

But as a function of the unnormalized filter, the volatility in the FB-HMM does not satisfy the assumptions in Haussmann and Sass (2004). In this work, we state the filtering equations for the FB-HMM, following the methodology from Elliott (1993), and see that the filter is as tractable as in the HMM from Haussmann and Sass (2004). Further, the issue of how the model behaves under time discretization is very important for practical applications, as we have already mentioned. For the question of approximation and consistency in the FB-HMM we consider the discretized returns. Here, the issue of consistency is a different one than what is typically investigated in statistics. The discretized FB-HMM depends on the continuous-time filter at the discretization points through its volatility, so it is not truly a discrete-time model. For consistency we consider the convergence of the discretized returns instead and prove that already for the simple Euler discretization the global discretization error converges to 0 in L^2.

The paper is organized as follows: in Section 2, we introduce the HMM and MSM in continuous time, prove the observability of the Markov chain in the continuous-time MSM and consider portfolio optimization in regime-switching models. Section 3 introduces the FB-HMM and prove an approximation result for a fixed information filtration. We further motivate the model by considering social learning arguments and shortly investigate the stylized facts present in the returns. Section 4 proves that the Euler-discretization of the FB-HMM returns converges to the continuous-time returns. We conclude the paper by presenting some numerical results and discuss an extension to a jump diffusion model.

2. The HMM and MSM in continuous time

Consider a return process $(R_t)_{t \in [0, T]}$ with dynamics

$$\mathrm{d}R_t = \mu_t \mathrm{d}t + \sigma_0 \mathrm{d}W_t,$$

where the unobservable drift process $\mu_t = \mu(Y_t)$ jumps between d states following a process Y. The volatility $\sigma_0 \in \mathbb{R}_{>0}$ is constant, where we denote with $\mathbb{R}_{d>0}$ the strictly positive real numbers. $(W_t)_{t \in [0, T]}$ is a standard Brownian motion defined