Journal of Econometrics 187 (2015) 418–435

Contents lists available at ScienceDirect

Journal of Econometrics

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

Market-based estimation of stochastic volatility models[☆]

Yacine Aït-Sahalia[a,b,*], Dante Amengual[c], Elena Manresa[d]

[a] Department of Economics, Princeton University, United States
[b] NBER, United States
[c] Centro de Estudios Monetarios y Financieros, Casado del Alisal 5, Madrid, 28014, Spain
[d] MIT Sloan School of Management, 30 Memorial Dr, Cambridge, MA 02142, United States

ARTICLE INFO

Article history:
Available online 23 April 2015

JEL classification:
G12
C32
C58

Keywords:
HJM approach
Variance swaps
Maximum likelihood estimation

ABSTRACT

We propose a method for estimating stochastic volatility models by adapting the HJM approach to the case of volatility derivatives. We characterize restrictions that observed variance swap dynamics have to satisfy to prevent arbitrage opportunities. When the drift of variance swap rates are affine under the pricing measure, we obtain closed form expressions for those restrictions and formulas for forward variance curves. Using data on the S&P500 index and variance swap rates on different time to maturities, we find that linear mean-reverting one factor models provide inaccurate representation of the dynamics of the variance swap rates while two-factor models significantly outperform the former both in and out of sample.

1. Introduction

We propose in this paper a novel method for estimating stochastic volatility models by adapting the approach of Heath et al. (1992) (HJM) from bond pricing to stochastic volatility models, using volatility derivatives as the market inputs. The HJM approach is motivated by the fact that in liquid markets, such as many markets for options and variance derivatives, one can rely on the observed market prices of standard liquid securities, which are expected to be fair under standard competitive markets assumptions, in order to estimate a model and subsequently price and hedge other more complex or less liquid securities.

In the early option pricing literature, starting with the seminal contributions of Black and Scholes (1973) and Merton (1973), stock prices were assumed to follow univariate processes with constant volatility. The empirical observation that options' implied volatility changes through time, across strike prices and maturities, rapidly gave rise to models relaxing this assumption, such as the stochastic volatility model of Hull and White (1987), or deterministic time-varying models for implied volatilities such as Derman and Kani (1994), Dupire (1993) and Dupire (1994). Stein and Stein (1991) and Heston (1993), among others, developed stochastic volatility models in which innovations to volatility need not be uncorrelated with innovations to the price of the underlying asset, unlike Hull and White (1987). These models proved useful in terms of explaining some of the joint time-series behavior of stock and option prices.

Estimating stochastic volatility models poses substantial challenges, though. Most of the difficulties arise from the fact that, by nature, volatility dynamics are not entirely observable. Further, the presence of one or more additional state variables that drive the volatility of the underlying asset price confers less tractability to the model from an analytic perspective.

In order to overcome the fact that volatility is unobservable, market-based perspectives have been adopted in the stochastic volatility literature. These methods take as primitives a set of financial securities which are liquidly traded. The basic assumption is that the fair price of each security is observable, and any quantity of the security can be sold or bought at its observed price at any point in time. These prices encapsulate what the market is telling the modeler. The market-based modeling approach postulates dynamic equations for the prices of the liquid instruments. The difficulty inherent in the approach consists in checking that the multitude of equations does not introduce inconsistencies and spurious arbitrage opportunities in the model: in other words, if

☆ We would like to thank participants at the Implied Volatility (Huntington Beach) and Measuring Risk Conferences (Princeton) for helpful comments and suggestions.

* Correspondence to: Bendheim Center for Finance, 26 Prospect Avenue, Princeton, NJ 08540, USA.

E-mail addresses: yacine@princeton.edu (Y. Aït-Sahalia), amengual@cemfi.es (D. Amengual), manresa@cemfi.es (E. Manresa).

http://dx.doi.org/10.1016/j.jeconom.2015.02.028
0304-4076/

the model is used to price the securities that served as its inputs, does it output back their prices?

Stochastic implied volatility models are an example of market-based models whose primitives are implied volatilities. This type of models focuses on the stochastic dynamics of either a single implied volatility (e.g. Lyons, 1997), the term structure of implied volatility (as in Schönbucher, 1999), or the whole volatility surface (e.g. Albanese et al. (1998); Ledoit et al. (2002)). While these approaches deal with the problem of stochastic implied volatility from a theoretical perspective, Rosenberg (2000) and Cont et al. (2002), among others, focus on the empirical aspects of the problem. However, estimating an arbitrage-free model of stochastic implied volatility is involved since the quantity of interest inherits the nonlinearities of the Black–Scholes formula (see e.g. Durrleman, 2004; Bergomi, 2005).

This paper contributes to the literature on stochastic volatility models by introducing the use of variance swap rates, which are the most actively traded variance related derivative products, instead of implied volatilities, as primitives for the stochastic volatility model, and then by deriving the constraints on the model to make it consistent with its inputs, detailing analytically the class of consistent models, and developing a closed-form likelihood-based estimation procedure for any consistent model.

A variance swap is a forward contract on the future quadratic variation that pays, at expiration, the difference between the quadratic variation over the horizon of the contract and the fixed variance swap rate. The use of variance swap rates offers several advantages for stochastic volatility modeling. In the first place, they are popular volatility derivatives, frequently used by market practitioners as a hedging tool for volatility risk, and have been actively traded in over-the-counter markets since the late 1990s, and in the simpler form of variance futures on exchanges. For that reason, variance swaps constitute a particularly reliable building block for a market based approach. Moreover, they can be easily synthesized from option prices (see e.g. Carr and Wu, 2009; Egloff et al., 2010; among others).[1] Second, for model tractability, they are especially convenient because of the linearity of their payoffs, together with the fact that variance derivatives can be seen as a family of conditional expectations over the distribution of the underlying's volatility, leading to linear relationships between the unobserved spot variance and the observed rates.

Of course, the model needs to be made consistent in the sense that if the model were used to price its inputs as any other derivative contracts, the model's output price should match the inputted price. We adopt a Markovian setup as in Bühler (2006) as our starting point. For an arbitrary number of sources of uncertainty in the economy, and assuming that we observe as many different time-to-maturity variance swap rates as sources of uncertainty, we characterize in analytical form the restrictions that (observed) derivative prices dynamics must satisfy in order to prevent arbitrage opportunities under the pricing measure. Moreover, we can implicitly define consistent forward variance curves. Consistent dynamic models are crucial for hedging purposes, in particular, where it is more important to fit the time series properties of the variables rather than the time to maturity or strike dimension of implied volatilities.[2]

When, additionally, (i) drift functions of variance swap rates are affine under the pricing measure Q, and (ii) forward variance curves are also affine in variance swap rates, we are able to obtain closed-form analytical expressions for the restrictions that affect the dynamics of observed prices and formulas for consistent forward variance curves. In particular, for a model with d sources of volatility risk, we find that a consistent affine–drift model is determined by $d + 1$ free parameters. These parameters can be naturally interpreted as speeds of mean reversion of the stochastic system and its steady state level.

Apart from its tractability, the affine-Q drift models have two additional attractive features worth mentioning. First, the diffusion functions of variance swap rates remain unrestricted except for the instantaneous variance of the underlying asset; hence, leaving room for non-affine specifications of the diffusion such as the constant elasticity of variance (CEV) model of Jones (2003), which is a relevant feature when fitting the models to the data; see, e.g., Li and Zhang (2013). Second, drift specifications richer than the ones belonging to the affine class can be obtained under the objective measure P through different specifications of market prices of risk.

For estimation, we rely on observations on the joint time-series of the underlying asset price and several variance swap rates with different maturities. Our joint modeling strategy allows us to separately pin down risk premia related to each of the different sources of uncertainty. However, as is usual in many financial applications dealing with multivariate diffusions, likelihood functions are not known explicitly for the models of interest. The solution to this problem relies on the approach of Aït-Sahalia (2008), who developed closed-form series expansions for the likelihood function for arbitrary multivariate continuous-time diffusions at discrete intervals of observations. Moreover, maximum likelihood estimation allows us to use likelihood ratio tests to evaluate the fit of non-nested models (see Vuong, 1989), which is difficult or impossible for other methods, such as the method of moments.

We derive explicit pricing formulas for two important classes of affine-Q drift models. The first one includes two sources of uncertainty in the economy, one specific to variance swap rates and one common to both variance swap rates and the stock; the second one assumes there are three sources of uncertainty, two of which driving the dynamics of variance swap rates. The first class of models we derive nests several stochastic volatility models already considered in the literature such as Heston (1993)'s, the GARCH stochastic volatility model considered in Nelson (1990), or the CEV model of Jones (2003). The second one generalizes stochastic volatility specifications with two volatility factors considered in Gatheral (2008), Amengual (2008), Christoffersen et al. (2009) and Egloff et al. (2010) among others, and is compatible with the evidence in favor of two volatility factors provided by Li and Zhang (2010).

We assess the small sample properties of the parameter estimates resulting from our proposed methodology by means of Monte Carlo experiments. Given the numerical tractability of the model and estimation method, we can easily perform large numbers of Monte Carlo simulations. In particular, we examine different scenarios including the effect of using different information sets, sampling frequency, near unit-root behavior for volatility and compare the small sample behavior of the estimators to their predicted asymptotic one. As expected, the sampling variance of diffusion parameters estimates is smaller than the corresponding to drift parameters, specially for those entering only through the P measure. Finally, with additional variance swap rates, the precision of the Q-drift parameters estimates significantly increases.

When taking the models to the data, we find that observation of additional variance swap rates with different time-to-maturity rates lead to significant efficiency gains in the estimation of the drift parameters under the pricing measure. Moreover, in agreement with the recent option pricing literature, our empirical results suggest that linear mean-reverting one factor models

[1] In practice, they may be a better choice than option prices since deep in- or out-of-the-money options tend to be illiquid. Using variance swap prices avoids this problem because prices of individual options are aggregated, so the impact of liquidity-induced idiosyncratic errors in each option price is reduced.

[2] The usual practice relies on calibrating alternative models using information of stock and option prices (equivalently implied volatilities) at a given point in time, which ultimately is not accounting for the dynamics of the processes.