

Appendix 2: Finite Sample Performance

In this appendix, we assess the small sample properties of the methodology proposed in Aït-Sahalia et al. (2014) through an extensive Monte Carlo design. Due to the analytical, fully closed-form, properties of the method, Monte Carlo simulations can be conducted at a relatively low computational cost. We analyze three different scenarios: The first one considers sample properties under different combinations of sample frequencies and sample sizes. The second is concerned with the accuracy of the estimations when the speed of mean reversion is relatively high and relatively low, to assess the performance of the method in finite samples when volatility has a near unit-root behavior (a feature documented in Aït-Sahalia (1996) for interest rates and Bollerslev et al. (2012) for volatility, among others). Finally, in the third one, we present results on efficiency gains when additional variance swap rates are used in the estimation relative to the exactly identified scenarios.

We expect the performance of the parameters to depend on how and where they appear in drift and diffusion functions, respectively. Notice that we can split the vector of parameters into three different blocks: parameters entering only the drift function, such as market prices of risk, coefficients entering the diffusion matrix and the P -drifts, as the parameters characterizing the correlation structure of the different sources of uncertainty ($\gamma's, \beta$), and parameters such as the long run mean ($\bar{\alpha}$), and the speeds of mean reversion ($\lambda's$), that are involved in the drift functions under both measures as well as in the diffusion matrix. The parameter δ is poorly identified; for that reason, we follow the standard approach in the option pricing literature square-root specification whereby δ remains fixed to 1. In any event, different values of δ have a relatively small effect in the estimation of the remaining parameters with the exception of ν 's which only adjust to correct the change of scale σ_t^δ .

We run Monte Carlo simulations for the 1-VS model presented in Section 2.4.1 since most of the relevant intuition can be extracted from this simple setup. We run simulations for a CEV specification of the diffusion function, which is a special case of the consistent 1-VS family of models. The CEV specification is a natural choice since it nests GARCH and Heston type stochastic volatility models that have been extensively studied in the empirical option pricing literature.

For each batch of simulations, we generate 1,000 sample paths using an Euler discretization of the process, with 25 sub-intervals per sampling interval; 24 out of 25 observations are then discarded. Each simulated data series is initialized with the volatility state variable at its unconditional mean. The first 100 generated observations are discarded, while the first of the n remaining observations is taken as the starting point for the simulated data series. Our benchmark sample has $n = 1,500$ observations at the daily frequency ($\Delta = 1/252$), which roughly corresponds to 6 years of daily data, on both the spot and the 2 months time-to-maturity variance swap rate. The speed of mean reversion is set to 3 and the unconditional level of volatility at 32% per year. As for the diffusion parameters and market prices of risk we fix them to standard values in the literature (see Table 6).

Results for different sampling frequencies and periods are reported in Table 6. We sample lengths of $n = 500$ and 4,000 transitions at the daily frequency ($\Delta = 1/252$) and $n = 250$ at the weekly frequency ($\Delta = 1/52$). These sample sizes correspond to 2, 16 and 5 years of data, respectively. The use of similar batches of simulations with differing numbers of observations in each simulated series provides

some insight into how fast the small-sample distribution of the estimated parameters approaches the asymptotic distribution. For a given sampling frequency, as the number of observations in each simulated data series increases, we would expect the standard errors of the parameter estimates to decrease at a rate inversely proportional to the square root of the number of observations. Indeed, the standard errors for all parameters decrease roughly by a factor of $\sqrt{3}$ when increasing the sample size from 500 to 1,500 observations. Throughout, the more precise estimates are obtained for the γ_1 , γ_{11} and β parameters. These patterns for γ 's and β are the same irrespective of the exercise we perform as can be also seen in Tables 7 and 8. The only exception corresponds to the case of weekly data, which is also in line with predictions from asymptotic theory for discretely sampled diffusions. By contrast, regardless of the sampling frequency and the sample length, both the biases and standard errors of the market prices of risks estimates are relatively large. As for λ_1 and $\bar{\alpha}$, they are estimated with less accuracy than the other coefficients entering into the diffusion matrix. Still, the bias in these estimates decreases substantially with the length of the sample, which is possibly indicative that the source of the bias identified in Li et al. (2004) could be relevant for sample sizes of 500 daily observations.

Next, we pay attention to the precision of the estimates of the parameters λ_1 and $\bar{\alpha}$ under DGPs in which we allow for different speed of mean reversion for the reference variance swap rate. Table 7 contains biases and standard deviations when λ_1 takes the values $-1/3$, -1 , -3 and -9 . The first two columns refer to the DGP with $\lambda_1 = -1/3$. In this case, both the speed of mean reversion λ_1 and the unconditional level $\bar{\alpha}$ are overestimated (i.e. the bias on the λ_1 parameter is positive) and the standard deviation of the empirical distribution of those parameter estimates is large. A potential explanation for the biases can be found in the fact that given the parametrization, we are only using six years of data to estimate a model in which the mean life of shocks to volatility is higher than two years. We also overestimate in small-samples the speed of mean reversion when it is either -1 or -3 , but with a lower bias than in the first case. Similarly, the bias of $\bar{\alpha}$ decreases monotonically as the speed of mean reversion of the process increases. Finally, for the fast mean-reverting configuration ($\lambda_1 = -9$) the estimates tend to underestimate λ_1 .

We finally analyze how increasing the information set used for inference affects the estimator's efficiency; we do so by including additional variance swap rates subject to pricing errors into the likelihood function. We first include the 6-month time-to-maturity rate to the benchmark case, then the 3-month, 6-month and 1-year rates, and finally also include 1-month and 2-year swap rates. We generate additional variance swap rates using the pricing relationships implied by Proposition 2 and we add i.i.d. multiplicative pricing errors with standard deviation equal to 0.1. Biases and standard deviations are in Table 8. Not surprisingly, efficiency gains are more noticeable for parameters that enter into the forward variance curve; but also these gains in precision help to reduce the standard deviations of the remaining parameters of the model, specially the market prices of risks ν 's. More generally, it is interesting to note that the inclusion of additional variance swap rates helps to better estimate the parameters, but once these gains have been achieved with a few additional rates, there is no big difference between using a larger set of rates: there are only minor differences in biases and standard deviations when using five instead of three additional swap rates with different time-to-maturities.

Daily observations ($\Delta = 1/252$)					
<i>Benchmark Case</i>					
$n = 1,500$					
True value	$n = 1,500$		$n = 500$		Std. dev.
	Bias	Std. dev.	Bias	Std. dev.	
Market-based model parameters					
$\bar{\alpha}$	0.1	-0.002	0.012	-0.012	0.069
λ_1	-3	0.081	0.643	0.390	1.305
ν_0	0.05	0.003	0.595	0.067	1.563
ν_1	-0.1	0.027	1.204	0.153	2.891
γ_1	-0.5	0.001	0.018	-0.001	0.032
γ_{11}	0.625	0.008	0.070	0.017	0.135
β	1.5	0.015	0.090	0.044	0.171
Daily observations ($\Delta = 1/252$)					
Weekly observations ($\Delta = 1/52$)					
$n = 4,000$					
True value	$n = 4,000$		$n = 250$		Std. dev.
	Bias	Std. dev.	Bias	Std. dev.	
Market-based model parameters					
$\bar{\alpha}$	0.1	-0.001	0.032	-0.029	0.180
λ_1	-3	-0.048	0.398	-0.317	1.036
ν_0	0.05	0.076	0.715	0.195	1.434
ν_1	-0.1	0.072	1.170	0.253	2.467
γ_1	-0.5	-0.000	0.011	-0.002	0.044
γ_{11}	0.625	-0.001	0.040	0.048	0.161
β	1.5	0.001	0.050	0.100	0.219

Table 6. Monte Carlo Simulations: 1-VS CEV-type Model with Different Sample Sizes and Sampling Frequencies.

Notes: This table shows the results of 1,000 Monte Carlo simulations, with different sample sizes and sampling frequencies, for the CEV model of the 2-months time-to-maturity variance swap rate using an Euler discretization of the process of 25 sub-intervals per sampling interval. The benchmark consists in 6 years of data, that is $n = 1,500$ observations, sampled at the daily frequency ($\Delta = 1/252$), while the remaining columns include sample lengths of $n = 500$ and 4,000 transitions at the daily frequency that correspond to 2 and 16 years of data. The last two columns correspond to 5 years of weekly data i.e. $n = 250$ transitions with $\Delta = 1/52$. The second column shows the true value θ_0 of the parameters used to generate the simulated sample paths. The “Bias” column shows the mean bias of the estimated parameter vector, i.e., the difference between the estimated parameters and the true values. The “Std. dev.” column shows the standard deviation of the parameter estimates. The models and parameterizations are given in Section 3.

	True value	$\lambda = -1/3$		$\lambda = -1$	
		Bias	Std. dev.	Bias	Std. dev.
Market-based model parameters					
$\bar{\alpha}$	0.1	-0.098	0.547	-0.030	0.160
λ_1		0.101	0.305	0.115	0.488
ν_0	0.05	0.023	0.530	0.065	0.730
ν_1	-0.1	0.086	1.241	0.069	1.094
γ_1	-0.5	-0.000	0.019	0.000	0.018
γ_{11}	0.625	0.009	0.045	0.008	0.056
β	1.5	0.012	0.053	0.012	0.066
<i>Benchmark Case</i>					
	True value	$\lambda = -3$		$\lambda = -9$	
		Bias	Std. dev.	Bias	Std. dev.
Market-based model parameters					
$\bar{\alpha}$	0.1	-0.002	0.012	-0.001	0.004
λ_1		0.081	0.643	-0.557	0.873
ν_0	0.05	0.003	0.595	0.123	1.038
ν_1	-0.1	0.027	1.204	0.095	1.555
γ_1	-0.5	0.001	0.018	-0.000	0.019
γ_{11}	0.625	0.008	0.070	-0.013	0.082
β	1.5	0.015	0.090	-0.008	0.102

Table 7. Monte Carlo Simulations: 1-VS CEV-type Model with Different Speed of Mean Reversion.

Notes: This table shows the results of 1,000 Monte Carlo simulations, with different speed of mean reversion λ , for the CEV model of the 2-months time-to-maturity variance swap rate using an Euler discretization of the process of 25 sub-intervals per sampling interval. The benchmark case employed consists of 6 years of daily data, that is $n = 1,500$ observations sampled at the daily ($\Delta = 1/252$) frequency, while the remaining columns include, for the same sample characteristics, different values of the mean reversion parameter λ of the variance swap rate. The second column shows the true value θ_0 of the parameters used to generate the simulated sample paths. The “Bias” column shows the mean bias of the estimated parameter vector, i.e., the difference between the estimated parameters and the true values. The “Std. dev.” column shows the standard deviation of the parameter estimates. The models and parameterizations are given in Section 3.

	True value	Benchmark Case		Additional VS rates	
		Bias	Std. dev.	6 months	
				Bias	Std. dev.
Market-based model parameters					
$\bar{\alpha}$	0.1	-0.002	0.012	0.000	0.001
λ_1	-3	0.081	0.643	0.004	0.096
ν_0	0.05	0.003	0.595	0.003	0.089
ν_1	-0.1	0.027	1.204	0.017	0.208
γ_1	-0.5	0.001	0.018	0.000	0.017
γ_{11}	0.625	0.008	0.070	0.004	0.048
β	1.5	0.015	0.090	0.007	0.060
Pricing errors for additional variance swap rates					
σ_{6M}	0.1			-0.0001	0.0018
Additional VS rates					
	True value	1, 6 & 12 m.		1,3, 6, 12 & 24 m.	
		Bias	Std. dev.	Bias	Std. dev.
Market-based model parameters					
$\bar{\alpha}$	0.1	0.000	0.000	0.000	0.000
λ_1	-3	0.000	0.054	0.004	0.046
ν_0	0.05	-0.003	0.025	-0.004	0.020
ν_1	-0.1	0.009	0.083	0.002	0.030
γ_1	-0.5	-0.001	0.018	0.000	0.017
γ_{11}	0.625	0.003	0.044	0.004	0.042
β	1.5	0.006	0.055	0.008	0.053
Pricing errors for additional variance swap rates					
σ_{1M}	0.1	-0.0001	0.0018	0.0000	0.0019
σ_{3M}	0.1			-0.0000	0.0019
σ_{6M}	0.1	-0.0000	0.0019	-0.0000	0.0018
σ_{1Y}	0.1	-0.0000	0.0019	-0.0001	0.0019
σ_{2Y}	0.1			0.0001	0.0018

Table 8. Monte Carlo Simulations: 1-VS CEV-type Model with Observation of Additional Variance Swap Rates.

Notes: This table shows the results of 1,000 Monte Carlo simulations, with observation of additional variance swap rates, for the CEV model with 2-months time-to-maturity variance swap rate using an Euler discretization of the process of 25 sub-intervals per sampling interval. The benchmark case employed consists in 6 years of daily data, that is $n = 1,500$ observations sampled at the daily ($\Delta = 1/252$) frequency, while the remaining columns include summary statistics for parameter estimates when additional variance swap rates are observed: different columns report results for alternative configurations including subsets of 1, 3, 6, 12 and 24 months time-to-maturity. The second column shows the true value θ_0 of the parameters used to generate the simulated sample paths. The “Bias” column shows the mean bias of the estimated parameter vector, i.e., the difference between the estimated parameters and the true values. The “Std. dev.” column shows the standard deviation of the parameter estimates. The models and parameterizations are given in Section 3.