Innovations in Insurance, Risk- and Asset Management

Kathrin Glau
Daniël Linders
Aleksey Min
Matthias Scherer
Lorenz Schneider
Rudi Zagst

editors

KPMG Center of Excellence
in Risk Management

World Scientific

Chapter 12

Option Pricing and Hedging for Discrete Time Autoregressive Hidden Markov Model

Massimo Caccia* and Bruno Rémillard†

*Department of Decision Sciences, HEC Montréal,
Montréal (Québec), Canada H3T 2A7*

*massimo.caccia@hec.ca

†bruno.remillard@hec.ca

In this paper we solve the discrete time mean-variance hedging problem when asset returns follow a multivariate autoregressive hidden Markov model. Time dependent volatility and serial dependence are well established properties of financial time series and our model covers both. To illustrate the relevance of our recommended approach, we first compare the proposed model with the well-known hidden Markov model via likelihood ratio tests and a novel goodness-of-fit test on the S&P 500 daily returns. In addition, we present out-of-sample hedging results on S&P 500 vanilla options as well as a trading strategy based on the difference between theoretical and market prices. This strategy is compared to simpler models including the classical Black-Scholes delta-hedging approach.

Keywords: option pricing, dynamic hedging, regime-switching, goodness-of-fit, auto-regressive hidden Markov models.

1. Introduction

The quest for the perfect option pricing model is an important topic in the mathematical finance literature. [1] provided the following observation: if a claim is priced by arbitrage in a world with one asset and one bond, then its value can be found by first adapting the model so that the asset earns the risk-free rate, and then computing the expected value of the claim. The idea of finding a self-financing optimal investment strategy that replicates the terminal payoff of the claim is now known as dynamic hedging.

One can model the underlying asset's returns with the geometric Brownian motion and retrieve a tractable and intuitive way of pricing and replicating options. This is precisely what [2] proposed; it is now called the B&S model. Unfortunately, financial markets are far too complex for a model as simple as this one and this hedging protocol can lead to large hedging errors, as it will be shown later in this paper. The main drawback of this

Open Access chapter published by World Scientific Publishing Company and distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives (CC BY-NC 4.0) License.

framework is the constant volatility assumption. Indeed, according to [3], [4], volatility seems to vary over time, mainly for macroeconomic reasons. Furthermore, the B&S model assumes serial independence for the returns, which is also an hypothesis that is violated in general.

Optimal hedging was introduced later, and it consists in minimizing the quadratic error of replication. Solutions were derived in continuous time in [5] and later in discrete time in [6]. This methodology can be applied to the B&S model, and more interestingly, to stochastic volatility models.

Hidden Markov models, popularized in [7], [8] were proven to be extremely useful for modeling economic and financial time series. They are robust to time-varying volatility, serial correlation and higher-order moments, which are all well-established stylized facts of asset returns. The premise for these models is that identifiable events can quickly change the characteristics of an asset's returns. This should be taken into account when pricing a derivative. These events could be on a long horizon — fundamental changes in monetary, fiscal or income policies — or on a shorter horizon — news related to the underlying stock or changes in the target band for the federal funds rate. However, the classical implementation of an HMM can't account for multiple horizons.

Elliot's work on energy finance and interest rate modeling, where mean-reversion is a widely accepted feature, addresses this problem. [9] introduced a way to parameterize a regime-switching mean-reverting model with jumps. They found the calibration of the model to be difficult because of the small amounts of jumps exhibited in the time series. [10] later introduced a similar model with no jumps, and where the volatility is subject to mean-reverting regime-switches. The basis was that volatility, being driven by macroeconomic forces, was not to be modeled by price movements. Hence the need to model it by a hidden Markov chain. Finally, [11] investigated the valuation of European and American options under another model where the volatility is subject to regime-switches, but this time the Markov chain was assumed to be observable. This suggests that it would be interesting to develop some methods to determine the optimal number of states for the hidden Markov chain. This is precisely one of the contributions of our paper.

In light of all the above, we decided to generalize the work of [12]: we combine the regime-switching model with an autoregressive parameter to account for trends and mean-reversions without having to change regime. Autoregressive hidden Markov models (ARHMM) have been applied to financial engineering and have shown promising results, see, e.g. [13]. Still,

this model has never been used in conjunction with optimal hedging. We derive the solution of the hedging strategy and obtain derivatives prices under this class of models. It is also noteworthy to add that we will use semi-exact techniques to compute expectations necessary for the optimal hedging, instead of Monte Carlo techniques, which will greatly speed up computations. For parameterization, we will implement the EM algorithm of [14] to the ARHMM. This method is widely used in unsupervised machine learning in order to find hidden structures, in our case, the regimes. In order to choose the optimal number of regimes and to assess the suitability of the model, we propose a new goodness-of-fit test based on [15] and [16]. It uses the Rosenblatt transform and parametric bootstrap. Compared to Elliot's work, our model can exhibit mean-reversion but is not restricted to it. It could thus be adequate for the modeling of a wide variety of assets.

In his famous study [17], Fama presented strong and voluminous evidence in favor of the random walk hypothesis. He although suggested that other tests — statistical or profit generating strategies — could either confirm or contradict his findings. In this paper, we will explore both avenues. We will statistically show that the ARHMM is an adequate model for financial modeling using the goodness-of-fit test as well as likelihood ratio tests, and we will show that it is possible to generate money by buying/selling options and replicating them until maturity. To support our approach, we will compare the trading strategy's returns with different methodologies: Black-Scholes delta-hedging and optimal hedging when assets follow a geometric random walk. We will also compare the hedging results with the delta-hedging using the market's implicit volatility.

First, we present likelihood ratio test results confirming the ARHMM is a better fit than the classical HMM on S&P daily returns, in particular, because our model has the capacity for mean-reversion. Secondly, empirical pricing and hedging results suggest that our methodology is superior to its counterparts by achieving the best mean-squared error in six out of eight cases, as well as being the most profitable strategy.

The rest of the paper is organized as follows. Section 2 describes the model and implements the EM algorithm for parameter estimation. In addition, we will introduce the goodness-of-fit test and study its suitability in the financial markets. Then, in Section 3, we will state the optimal dynamic discrete time hedging model when assets follow a ARHMM. The results of the implementation of the dynamic hedging strategies will be presented in Section 4. Section 5 concludes.

2. Regime-switching autoregressive models

The proposed models are quite intuitive. The regime process τ is a homogeneous Markov chain on $\{1, \dots, l\}$, with transition matrix Q . At period $t - 1$, if $\tau_{t-1} = i$, and the return Y_{t-1} has value y_{t-1} , then at time t , $\tau_t = j$ with probability Q_{ij} , and the return Y_t has conditional distribution $f_j(y_t; y_{t-1})$; here lower case letters $y_1, \dots, y_n$ are used to denote a realization of $Y_1, \dots, Y_n$. From this construction it follows that (Y_t, τ_t) is a Markov process. For example, for $j \in \{1, \dots, l\}$, one could take a Gaussian AR(1) model, meaning that given $Y_{t-1} = y_{t-1}$ and $\tau_t = j$, $Y_t = \mu_j + \Phi_j(y_{t-1} - \mu_j) + \varepsilon_t$, with $\varepsilon_t \sim N(0, A_j)$; the conditional density of Y_t at $y_t \in \mathbb{R}^d$ is

$$f_j(y_t|y_{t-1}) = \frac{e^{-\frac{1}{2}\{y_t - \mu_j - \Phi_j(y_{t-1} - \mu_j)\}^\top A_j^{-1}\{y_t - \mu_j - \Phi_j(y_{t-1} - \mu_j)\}}}{(2\pi)^{d/2}|A_j|^{1/2}}, \quad (1)$$

where $\mu_j \in \mathbb{R}^d$, Φ_j is a $d \times d$ matrix such that $\Phi_j^n \rightarrow 0$ as $n \rightarrow \infty$,¹ and A_j is a $d \times d$ non-degenerate covariance matrix. The matrices $\Phi_1, \dots, \Phi_l$ are mean-reversion parameters. Let $\mathcal{B}_d$ be the set of $d \times d$ matrices B such that $B^n \rightarrow 0$ as $n \rightarrow \infty$ and let S_d^+ be the set of symmetric positive definite $d \times d$ matrices. Note that $\mathcal{B}_d$ is the set of $d \times d$ matrices with spectral radius smaller than 1, meaning that the eigenvalues are all in the unit complex ball of radius 1; in particular, $I - B$ is invertible for any $B \in \mathcal{B}_d$. Note that the so-called Hidden Markov Model is obtained by setting $\Phi_1 = \dots = \Phi_l = 0$.

2.1. Regime prediction

Since the regimes are not observable, we have to find a way to predict them. This will be of utmost importance for pricing and hedging derivatives.

In many applications, one has to predict a colored a non-observable signal by using observations $Y_1, \dots, Y_t$ linked in a certain way to the signal. This is known as a filtering problem; see, e.g., [18]. In our case, we need to find at time t the conditional probability $\eta_t(i) = P(\tau_t = i | Y_1 = y_1, \dots, Y_t = y_t)$. It is remarkable that for the present model, one can compute exactly this conditional distribution, given a starting distribution η_0 . For more details, see the extension of the Baum-Welch algorithm in Appendix A.

¹This condition ensures that for any $j \in \{1, \dots, l\}$, the matrix $B_j = \sum_{k=0}^{\infty} \Phi_j^k A_j (\Phi_j^k)^\top$ is well defined and satisfies $B_j = \Phi_j B_j \Phi_j^\top + A_j$.

2.1.1. Filtering algorithm

Choose an a priori probability distribution η_0 for the regimes. Equivalently, one can choose a positive vector q_0 and set $\eta_0(i) = q_0(i)/Z_0$, where $Z_0 = \sum_{j=1}^l q_0(j)$. The choice of q_0 or η_0 is not critical since its impact on predictions decays in time and have virtually no impact on terminal regime probabilities for any reasonable time series length. For simplicity, we assume a uniform distribution, i.e. $q_0 \equiv 1/l$.

For $t \in \{1, \dots, n\}$, define $q_t(i) = E \left[\mathbf{1}(\tau_t = i) \prod_{k=1}^t f_{\tau_k}(y_k|y_{k-1}) \right]$, $i \in \{1, \dots, l\}$, and $Z_t = \sum_{j=1}^l q_t(j)$. Note that the first observation of the sequence is burn-in in order to compute $f_{\tau_1}(y_1|y_0)$. Hence, Z_t is the joint density $f_{1:t}(y_1, \dots, y_t)$ of $Y_1, \dots, Y_t$ because

$$f_{1:t}(y_1, \dots, y_t) = E \left[\prod_{k=1}^t f_{\tau_k}(y_k|y_{k-1}) \right] = \sum_{j=1}^l q_t(j) = Z_t.$$

Then if $q_0 = \eta_0$, for any $i \in \{1, \dots, l\}$, and any $t \geq 1$, we have

$$\begin{aligned} q_t(i) &= E \left[\mathbf{1}(\tau_t = i) \prod_{k=1}^t f_{\tau_k}(y_k|y_{k-1}) \right] \\ &= f_i(y_t|y_{t-1}) \sum_{j=1}^l E \left[\mathbf{1}(\tau_t = i) \mathbf{1}(\tau_{t-1} = j) \prod_{k=1}^{t-1} f_{\tau_k}(y_k|y_{k-1}) \right] \\ &= f_i(y_t|y_{t-1}) \sum_{j=1}^l Q_{ji} E \left[\mathbf{1}(\tau_{t-1} = j) \prod_{k=1}^{t-1} f_{\tau_k}(y_k|y_{k-1}) \right] \\ &= f_i(y_t|y_{t-1}) \sum_{j=1}^l Q_{ji} q_{t-1}(j), \end{aligned} \quad (2)$$

and

$$\eta_t(i) = P(\tau_t = i | Y_1, \dots, Y_t) = \frac{q_t(i)}{Z_t}. \quad (3)$$

Having computed the conditional distribution η_t , one can estimate τ_t by

$$\tau_t = \arg \max_i \eta_t(i), \quad (4)$$

i.e. as the most probable regime. In view of applications, it is preferable to rewrite (3) only in terms of η , i.e.,

$$\eta_t(i) = \frac{f_i(y_t|y_{t-1})}{Z_{t|t-1}} \sum_{j=1}^l \eta_{t-1}(j) Q_{ji},$$

where $Z_{t|t-1} = \frac{Z_t}{Z_{t-1}} = \sum_{j=1}^l \sum_{i=1}^l f_i(y_t|y_{t-1})\eta_{t-1}(j)Q_{ji}$. As a result, $Z_{t|t-1}$ is the conditional density of Y_t at y_t , given $Y_1 = y_1, \dots, Y_{t-1} = y_{t-1}$.

2.1.2. Conditional distribution

From the results of the previous section, the joint density $f_{1:t}$ of $Y_1, \dots, Y_t$ is Z_t . Also, for any $t \geq 2$, the conditional density $f_{t|t-1} = Z_{t|t-1}$ of Y_t given $Y_1, \dots, Y_{t-1}$, can be expressed as a mixture, viz.

$$\begin{aligned} f_{t|t-1}(y_t|y_1, \dots, y_{t-1}) &= \sum_{i=1}^l f_i(y_t|y_{t-1}) \sum_{j=1}^l \eta_{t-1}(j)Q_{ji} \\ &= \sum_{i=1}^l f_i(y_t|y_{t-1})W_{t-1}(i) \end{aligned}$$

where $W_{t-1}(i) = \sum_{j=1}^l \eta_{t-1}(j)Q_{ji}$, $i \in \{1, \dots, l\}$. Note that for all $t > 1$, $W_{t-1}(i) = P(\tau_t = i | Y_{t-1} = y_{t-1}, \dots, Y_1 = y_1)$. As a result, it follows that

$$P(\tau_{t+k} = i | Y_t = y_t, \dots, Y_1 = y_1) = \sum_{j=1}^l (Q^k)_{ji} \eta_t(j), \quad i \in \{1, \dots, l\}.$$

Next, the conditional law of $Y_{t+1}, \dots, Y_{t+m}$ given $Y_1, \dots, Y_t$ has density

$$\begin{aligned} f_{t+m|t}(y_{t+1}, \dots, y_{t+m} | y_1, \dots, y_t) &= \sum_{i_0=1}^l \sum_{i_1=1}^l \cdots \sum_{i_m=1}^l \eta_t(i_0) \\ &\quad \times \prod_{k=1}^m Q_{i_{k-1}i_k} f_{i_k}(y_{t+k} | y_{t+k-1}). \end{aligned}$$

2.1.3. Stationary distribution in the Gaussian case

Suppose that the model specified by (1) holds, ergo the innovations are Gaussian. If Y_n converges in law to a stationary distribution for any given starting point y_0 , then this distribution must be Gaussian, with mean μ and covariance matrix A . Suppose the Markov chain is ergodic with stationary distribution ν . Then with probability ν_i , $i \in \{1, \dots, l\}$, $Y_1 = (I - \Phi_i)\mu_i + \Phi_i Y_0 + \epsilon_i$, where $\epsilon_i \sim N(0, A_i)$ is independent of $Y_0 \sim N(\mu, A)$. It then follows that

$$\mu = \left\{ \sum_{i=1}^l \nu_i (I - \Phi_i) \right\}^{-1} \left\{ \sum_{i=1}^l \nu_i \{ (I - \Phi_i) \mu_i \} \right\}.$$

Similarly, A must satisfies $A = T(A)$, where

$$T(A) = B + \sum_{i=1}^l \nu_i \Phi_i A \Phi_i^\top, \quad (5)$$

with $B = -\mu\mu^\top + \sum_{i=1}^l \nu_i [(I - \Phi_i)\mu_i + \Phi_i\mu]\{(I - \Phi_i)\mu_i + \Phi_i\mu\}^\top + A_i]$. From the conditions on $\Phi_1, \dots, \Phi_l$, there is a norm $\|\cdot\|$ on the space of matrices such that $\|\Phi_i\| < 1$ for every $i \in \{1, \dots, l\}$.² The operator T is then a contraction since for any two matrices A_0, A_1 , $\|T(A_1) - T(A_0)\| \leq \|A_1 - A_0\| \sum_{i=1}^l \nu_i \|\Phi_i\|^2 \leq c \|A_1 - A_0\|$, with $c = \max_{1 \leq i \leq l} \|\Phi_i\|^2 < 1$. Also, since $T(A)$ is a covariance matrix whenever A is one, and B is positive definite, it follows that there is a unique fixed point A of T , meaning that $A = T(A)$, and this unique fixed point A is a positive definite covariance matrix. In fact, A is the limit of any sequence $A_n = T(A_{n-1})$, with A_0 a non-negative definite covariance matrix. For example, one could take even take $A_0 = 0$. This provides a way to approximate the limiting covariance A by setting $A \approx A_n$ for n large enough.

2.2. Estimation of parameters

The EM algorithm of [14] is an efficient estimation procedure for incomplete datasets. This is the case here since τ is unobservable. The algorithm proceeds iteratively to converge to the maximum likelihood estimation of parameters. Its implementation for ARHMM is detailed in Appendix A.1. It seems that starting the parameter's estimation of the ARHMM with the HMM parameters' estimate (obtained by setting $\Phi_1 = \dots = \Phi_d = 0$) is slightly more stable. The optimal number of regimes must be known a priori, an issue we will discuss next.

2.3. Goodness-of-fit test and selection of the number of regimes

To select the optimal number of regimes, one must test the adequacy of fitted models with different number of regimes. This is generally done by using a test based on likelihoods. However, according to [19], goodness-of-fit tests based on likelihoods are not recommended for regime-switching models. We opt for a simpler approach based on a parametric bootstrapping. It was shown to work on a large number of dynamic models, including hidden

²Recall that all norms are equivalent.

Table 1. P-values (in percentage) for the nonparametric change point test using the Kolmogorov-Smirnov statistic with N=10000 bootstrap samples.

Period	<i>P</i> -value
2000's recovery	39.8
2008-2009 Financial Crisis	0.4
2010's recovery	9.7

Markov models. The test is built on the work of [20] and its implementation is given in Appendix B.

2.3.1. *Selecting the number of regimes*

The goodness-of-fit test described in Appendix B produces P-values from Cramér-von Mises type statistics, for a given number of regimes ℓ . As suggested in [21], it makes sense to choose the optimal number of regimes ℓ^* as the smallest ℓ for which the P-value is larger than 5%. An illustration of this methodology is given in Section 2.4.

2.4. *Application to S&P 500 daily returns*

To assess the relevance of our model on real data, we estimated the parameters on the close-to-close log-returns of the daily price series of the S&P 500 Total Return. To find stationary estimation windows, we used a nonparametric changepoint test for a univariate series using a Kolmogorov-Smirnov type statistic, as in [18]. We focused on recent data, i.e. from early 2000 to the beginning of 2017. We found two stationary estimation windows: from 05/01/2004 to 02/01/2008 and from 05/01/2010 to 20/01/2017. We refer to the former as the 2000's recovery and 2010's recovery for the latter. Results of the tests are presented in Table 1. We also studied the interesting period in between, the 2008-2009 Financial Crisis, even though the null hypothesis of stationarity has a *P*-value of 0.4%.

Next, we performed the goodness-of-fit test (GoF for short) described in Appendix B for the ARHMM (AR(1)) as well as for the HMM (AR(0)), as a mean of comparison. The results are presented in Tables 2, 4 and 6. According to the selection method described in Section 2.3.1, we selected a three-regime model for the 2000's recovery, since 3 is the smallest number of regimes for which the *P*-value is larger than 5%. This is also true for the HMM model. Likewise, we chose a three-regime model for the 2008-2009 Financial Crisis, and a four-regime model for the 2010's recovery. Note that in the case of the 2010's Bull markets, a four regime model for the HMM

was not enough to get a P -value $> 5\%$. Furthermore, to measure the significance of ARHMM over HMM, we performed a likelihood ratio test. This is possible because the HMM is a special case of the ARHMM corresponding to $\Phi_1 = \dots = \Phi_\ell = 0$. The corresponding statistic is computed as follows:

$$D = -2 \log \left(\frac{L(\hat{\theta}_0|x)}{L(\hat{\theta}_1|x)} \right) = -2 \log \left(\frac{f_{1:n}(y_1, \dots, y_n|\hat{\theta}_0)}{f_{1:n}(y_1, \dots, y_n|\hat{\theta}_1)} \right)$$

where $\hat{\theta}_0$ are the model's parameters estimated under the null hypothesis, i.e. $\Phi_1 = \dots = \Phi_\ell = 0$, so the returns follow a Gaussian hidden Markov model, and $\hat{\theta}_1$ are the model's parameters estimated under the alternative, i.e. returns follow an autoregressive hidden Markov model. Under the null hypothesis, the limiting distribution of this statistic is a chi-square distribution with ℓ degrees of freedom, which is the number of extra parameters in the alternative model. Hence, under the null hypothesis, $D \sim \chi^2(\ell)$. The log-likelihoods of both models, the statistical test D and the χ^2 critical value at a significance level of 5% are also presented in Tables 2, 4 and 6. We clearly reject the null hypothesis for all models, proving we should favor ARHMM over HMM for each dataset.

The estimated parameters for the three periods are presented in Tables 3, 5, and 7, where the mean and standard deviation of each AR(1) and AR(0) Gaussian regime density f_i are respectively denoted by μ_i and σ_i , and are presented as annualized percentages values. The tables further contain the stationary regime probabilities ν , together with the estimated transition matrix Q . Regimes are ordered by increasing volatility σ_i , and incidentally by decreasing mean μ_i .

Table 2. P-values (in percentage) for the proposed goodness-of-fit test using N=10000 bootstrap samples on the S&P 500 daily returns for the 2000's recovery, along with the log-likelihood of the models and the P-values (in percentage) of the likelihood ratio test statistic D.

	Number of regimes		
	1	2	3
GoF P -value (ARHMM)	0	0	26.51
GoF P -value (HMM)	0	0	25.12
Log-likelihood (ARHMM)	3479	3542	3559
Log-likelihood (HMM)	3475	3539	3552
P -value (D)	0.43	3.63	0.18

In the case of the 2000's recovery, regime 1 is associated to Bull markets, which are characterized by strong positive premium and low risk ($\mu_1 =$

Table 3. Parameters estimation for the three-regime models on the S&P 500 Total Return daily returns for the 2000's recovery.

Parameter	AR(0) Regime			AR(1) Regime		
	1	2	3	1	2	3
μ	31.41	13.88	-17.23	34.89	6.99	-21.60
σ	2.18	10.09	18.02	3.34	11.03	18.95
Φ	0	0	0	-0.14	0.03	-0.19
ν	0.11	0.65	0.24	0.19	0.63	0.18
Q	0	0.92	0.08	0	0.96	0.04
	0.17	0.83	0	0.32	0.68	0
	0	0.03	0.97	0	0.04	0.96

Note: μ and σ are presented as annualized percentage.

Fig. 1. Most probable regimes for the three-regime AR(1) model fitted on the S&P 500 Total Return index from 05/01/2004 to 02/01/2008 together with the cumulative performance of the index. Darker areas represent higher volatility states.

35.89 and $\sigma_i = 3.34$). It seems that this state is intermittent in the sense that the Markov chain does not stay or has a very small probability of staying in regime 1 since $Q_{11} \approx 0$. However, this state is not due to outliers since the percentage of time the Markov chain is in this state is 11% for the HMM and 19% for the ARHMM. Regime 2 is an intermediate state. Lastly, regime 3 is associated with bear markets or corrections, as highlighted by the negative premium of -21.60 and the volatility of 18.95. The regimes are less distinct in the HMM case. Also, the likelihood ratio test statistic $D = 15.04$ informs us that ARHMM is a much better fit for the data.

Fig. 2. Probability of the regimes, i.e. η_t , for the three-regime AR(1) model fitted on the S&P 500 Total Return index from 05/01/2004 to 02/01/2008.

Indeed, we observe strong mean-reversion in regime 1 and 3 ($\Phi_1 = -0.14$ and $\Phi_3 = -0.19$). Figure 1 displays the filtered most probable regimes (as defined in Section 2.1) for the whole time series. The regimes are depicted by different shades of grey, ranging from dark for the high volatility regime to white for the low volatility regime. The probabilities η_t of each regime are displayed in Figure 2. Interestingly enough, the crisis in the subprime mortgage market is adequately captured by the high risk regime.

Table 4. P-values (in percentage) for the proposed goodness-of-fit test using $N=10000$ bootstrap samples on the S&P 500 daily returns for the 2008-2009 Financial Crisis, along with the log-likelihood of the models and the P-values (in percentage) of the likelihood ratio test statistic D .

	Number of regimes		
	1	2	3
GoF P -value (ARHMM)	0	0	59.59
GoF P -value (HMM)	0	0	72
Log-likelihood (ARHMM)	1,209	1,323	1,334
Log-likelihood (HMM)	1,214	1,318	1,329
P -value (D)	0.19	0.95	2.59

The second period studied is quite interesting. For the 2008-2009 Financial Crisis, regimes are extremely polarized, with expected returns ranging

Table 5. Parameters estimation for the three-regime models on the S&P 500 Total Return daily returns for the 2008-2009 Financial Crisis.

Parameter	AR(0) Regime			AR(1) Regime		
	1	2	3	1	2	3
μ	74.97	-5.28	-66.87	72.22	-0.42	-64.73
σ	5.52	23.17	57.80	5.12	22.56	55.68
Φ	0	0	0	-0.03	-0.16	-0.15
ν	0.15	0.58	0.27	0.14	0.58	0.28
Q	0	0.97	0.03	0	0.98	0.02
	0.25	0.75	0	0.23	0.77	0
	0	0.01	0.99	0	0.01	0.99

Note: μ and σ are presented as annualized percentage.

Fig. 3. Most probable regimes for the three-regime AR(1) model fitted on the S&P 500 Total Return index from 03/01/2008 to 04/01/2010, together with the cumulative performance of the index. Darker areas represent higher volatility states.

from 72.22 to -64.73. The bear market regime, i.e. regime 3, is exceptionally persistent and volatile, as highlighted by $Q_{3,3} = 0.99$ and $\sigma_3 = 55.68$. Once more, we find two regimes exhibiting mean-reversion, i.e. $\Phi_2 = -0.16$ and $\Phi_3 = -0.15$. Figures 3 and 4 are analogous to Figures 1 and 2 respectively. We can see that the Markov chain switched to the high risk regime right after the collapse of Lehman Brothers. Remarkably, it stayed in that regime throughout almost all the Banking Crisis, even though we observe numerous small upwards trends, meaning many thought we hit the bottom.

Fig. 4. Probability of the regimes, i.e. η_t , for the three-regime AR(1) model fitted on the S&P 500 Total Return index from 03/01/2008 to 04/01/2010.

Table 6. P-values (in percentage) for the proposed goodness-of-fit test using N=10000 bootstrap samples on the S&P 500 daily returns for the 2010's recovery, along with the log-likelihood of the models and the P-values (in percentage) of the likelihood ratio test statistic D.

	Number of regimes			
	1	2	3	4
GoF P -value (ARHMM)	0	0	0	1.56
GoF P -value (HMM)	0	0	0	5.83
Log-likelihood (ARHMM)	5,696	5,936	5,985	6,012
Log-likelihood (HMM)	5,694	5,931	5,981	6,006
P -value (D)	4.19	1.52	4.46	3.02

Table 7. Parameters estimation for the four-regime models on the S&P 500 Total Return daily returns for the 2010's recovery.

Parameter	AR(0) Regime				AR(1) Regime			
	1	2	3	4	1	2	3	4
μ	29.41	303.07	-68.77	-28.25	32.04	365.83	-68.82	-27.50
σ	6.70	8.13	13.50	29.05	6.73	7.56	13.39	28.50
Φ	0	0	0	0	-0.04	0.14	-0.09	-0.08
ν	0.44	0.09	0.33	0.14	0.45	0.09	0.32	0.14
Q	0.91	0	0.09	0	0.90	0	0.10	0
	0.47	0.07	0.45	0	0.53	0.06	0.41	0
	0	0.23	0.76	0.01	0	0.23	0.76	0.01
	0	0.03	0	0.97	0	0.03	0	0.97

Note: μ and σ are presented as annualized percentage.

For the last period, we chose a model with four regimes. As noted previously, the four-regime HMM did not pass the goodness-of-fit test. We still present the estimated parameters in Table 7 as a mean of comparison.

Fig. 5. Most probable regimes for the four-regime AR(1) model fitted on the S&P 500 Total Return index from 05/01/2010 to 20/01/2017, together with the cumulative performance of the index. Darker areas represent higher volatility states.

Fig. 6. Probability of the regimes, i.e. η_t , for the four-regime AR(1) model fitted on the S&P 500 Total Return index from 05/01/2010 to 20/01/2017.

The calibration for this period is less intuitive than the previous ones. The inverse correlation between risk and expected premium is not as strong. Also, both models have a non-persistent regime with huge expected returns,

(i.e. regime 2). Nevertheless, we still find modest mean-reversion for two regimes (i.e. regimes 3 and 4), and the high-risk regime is highly persistent, as highlighted by $Q_{4,4} = 0.97$, as it was in the two previous cases. The most probable regimes are displayed in Figure 5, while the probabilities of each regime are presented in Figure 6. Interestingly enough, the crisis in the subprime mortgage market is adequately captured by the high risk regime. The final part of 2011 was marked by fear of the European sovereign debt crisis spreading to Italy and Spain. Once again, the ARHMM isolated the stock markets fall quite accurately. We also estimated the ARHMM on the returns from 01/04/1999 to 01/20/2017. This long period is far from stationary, but it is still interesting to see how the model performs through recessions and recoveries. We chose a four-regime model, as indicated by the goodness-of-fit tests. We can see on Figure 7 that the 2000's bubble burst and the recent financial meltdown (2008-2009) are both correctly captured by the high risk regimes.

Fig. 7. Most probable regimes for the four-regime AR(1) model fitted on the S&P 500 Total Return index from 01/04/1999 to 01/20/2017, together with the cumulative performance of the index. Darker areas represent higher volatility states.

3. Optimal discrete time hedging

In what follows, we use the notations and results from [22]. Denote the price process by S , i.e., S_t is the value of d underlying assets at period t and let $\mathbb{F} = \{\mathcal{F}_t, \quad t = 0, \dots, n\}$ a filtration under which S is adapted. Further

assume that S is square integrable. Set $\Delta_t = \beta_t S_t - \beta_{t-1} S_{t-1}$, where the discounting factors $\beta_t = e^{-r_t}$ are deterministic for $t = 1, \dots, n$. We are interested in the optimal initial investment amount V_0 and the optimal predictable investment strategy $\vec{\varphi} = (\varphi_t)_{t=1}^n$ that minimize the expected quadratic hedging error for a given payoff C at time n (e.g. a call option). Formally, the problem is stated as

$$\inf_{\{V_0, \vec{\varphi}\}} E[\{G(V_0, \vec{\varphi})\}^2], \quad (6)$$

where $G = G(V_0, \vec{\varphi}) = \beta_n(C - V_n)$ and V_t is the current value of the replicating portfolio at time t . In other words, $\beta_t V_t = V_0 + \sum_{j=1}^t \varphi_j^\top \Delta_j$, for $t = 0, \dots, n$. To solve (6), set $P_{n+1} = 1$, and define, for $t = n, \dots, 1$,

$$\begin{aligned} \gamma_{t+1} &= E(P_{t+1} | \mathcal{F}_t), \\ \mathbf{a}_t &= E(\Delta_t \Delta_t^\top P_{t+1} | \mathcal{F}_{t-1}) = E(\Delta_t \Delta_t^\top \gamma_{t+1} | \mathcal{F}_{t-1}), \\ \mathbf{b}_t &= E(\Delta_t P_{t+1} | \mathcal{F}_{t-1}) = E(\Delta_t \gamma_{t+1} | \mathcal{F}_{t-1}), \\ \rho_t &= \mathbf{a}_t^{-1} \mathbf{b}_t, \\ P_t &= \prod_{j=t}^n (1 - \rho_j^\top \Delta_j), \end{aligned}$$

We now state Theorem 1 of [22], which is a multivariate extension of [6].

Theorem 3.1. *Suppose that $E(P_t | \mathcal{F}_{t-1}) \neq 0$ P-a.s., for $1, \dots, n$. This condition is always respected for regime-switching models. Then, the solution $(V_0, \vec{\varphi})$ of the minimization problem (6) is $V_0 = E(\beta_n C P_1) / E(P_1)$, and*

$$\varphi_t = \alpha_t - \check{V}_{t-1} \rho_t, \quad t \in \{1, \dots, n\}. \quad (7)$$

where

$$\alpha_t = \mathbf{a}_t^{-1} E(\beta_n C \Delta_t P_{t+1} | \mathcal{F}_{t-1}). \quad (8)$$

and $\check{S}$ and $\check{V}$ are the present values of S and V .

Remark 3.1. V_0 is chosen such that the expected hedging error, G , is zero. [22] also showed that $C_t(S_t, \tau_t)$ given by

$$\beta_t C_t = \frac{E(\beta_n C P_{t+1} | \mathcal{F}_t)}{E(P_{t+1} | \mathcal{F}_t)} \quad (9)$$

is the optimal investment at period t so that the value of the portfolio at period n is as close as possible to C in terms of mean square error G , in particular, $V_0 = C_0$. C_t can be interpreted as the option price at period t . By increasing the number of hedging periods, C_t should tend to a price

under a risk-neutral measure; see, e.g., [23]. For example, when there is only one regime, the density is Gaussian and Φ_1 fixed at 0, C_t tends to the usual Black-Scholes price. The optimal hedging implementation for ARHMM is described in [24]. It follows that

$$\check{C}_{t-1} = \beta_{t-1} C_{t-1} \gamma_t = E\{(1 - \rho_t^\top \Delta_t) \check{C}_t | \mathcal{F}_{t-1}\}, \quad (10)$$

$$\alpha_t = \mathbf{a}_t^{-1} E(\check{C}_t \Delta_t | \mathcal{F}_{t-1}). \quad (11)$$

To derive the optimal hedging algorithm, we also need the following result, valid for a general ARHMM. First, write $S_t = D(S_{t-1})e^{Y_t}$, where e^{Y_t} is the vector with components $e^{(Y_t)_j}$, and $D(s)$ is the diagonal matrix with diagonal elements $(s)_j$, $j \in \{1, \dots, d\}$. The proof of the following theorem is given in [24].

Theorem 3.2. For any $t \in \{1, \dots, n\}$, set

$$\mathbf{a}_t = D(\check{S}_{t-1})a_t(Y_{t-1}, \tau_{t-1})D(\check{S}_{t-1}),$$

$$\mathbf{b}_t = D(\check{S}_{t-1})b_t(Y_{t-1}, \tau_{t-1}),$$

$$\rho_t = D^{-1}(\check{S}_{t-1})h_t(Y_{t-1}, \tau_{t-1}),$$

$$\gamma_t = g_t(Y_{t-1}, \tau_{t-1}),$$

with $h_t = a_t^{-1}b_t$, where a_t , b_t , and g_t are deterministic functions given by

$$a_t(y, i) = E\{\zeta_t \zeta_t^\top g_{t+1}(Y_t, \tau_t) | Y_{t-1} = y, \tau_{t-1} = i\}, \quad (12)$$

$$b_t(y, i) = E\{\zeta_t g_{t+1}(Y_t, \tau_t) | Y_{t-1} = y, \tau_{t-1} = i\}, \quad (13)$$

$$g_t(y, i) = E\{g_{t+1}(Y_t, \tau_t) | Y_{t-1} = y, \tau_{t-1} = i\} - b_t^\top(Y_{t-1}, \tau_{t-1})h_t(Y_{t-1}, \tau_{t-1}), \quad (14)$$

with $\zeta_t = e^{Y_t - r_t} - \mathbf{1}$, and $g_{n+1} \equiv 1$.

If in addition $\beta_n C = \Psi_n(\check{S}_n)$, then $\check{C}_t = \Psi_t(\check{S}_t, Y_t, \tau_t)$, where

$$\Psi_{t-1}(s, y, i)$$

$$= E[\Psi_t\{D(s)e^{Y_t - r_t}, Y_t, \tau_t\} \{1 - h_t(y, i)^\top \zeta_t\} | Y_{t-1} = y, \tau_{t-1} = i], \quad (15)$$

and

$$\alpha_t = D^{-1}(\check{S}_{t-1})a_t^{-1}(Y_{t-1}, \tau_{t-1})\mathbf{A}_t(\check{S}_{t-1}, Y_{t-1}, \tau_{t-1}), \quad (16)$$

where

$$\mathbf{A}_t(s, y, i) = E[\Psi_t\{D(s)e^{Y_t - r_t}, Y_t, \tau_t\} \zeta_t | Y_{t-1} = y, \tau_{t-1} = i]. \quad (17)$$

For example, for a call option with strike K , $\Psi_n(s) = \max(0, s - \beta_n K)$.

3.1. Implementation issues

There are two main problems related to the implementation of the hedging strategy: a_t , b_t , g_t , Ψ_t and A_t defined in expressions (12)-(17) must be approximated and regimes must be predicted. We approximate a_t , b_t , g_t for values y on a finite grid and then use linear interpolation. In a similar manner, we approximate Ψ_t and A_t for values s and y on a (product) grid and then use interpolation. This way, the recursion formulas given by (15)-(17) can be solved. This approach was proposed in [12], where stratified Monte Carlo sampling was also used. Because the simulations are computationally expensive and introduce variability, we propose a novel technique to approximate these integrals using semi-exact calculations, inspired by [18] Chapter 3. The details for the semi-exact calculations are presented in [24].

We also tested the Monte Carlo sampling procedure as a mean of comparison. Interestingly, we found that by simply rescaling the Monte Carlo samples to the desired mean and volatility, we achieved results as accurate as the semi-exact calculations, as pointed out in Section 3.3.

As for defining the points on the grids, previous literature suggest choosing 10^3 equidistant points marginally covering at least 3 standard deviations under the respective highest volatility regimes. Importantly, we found that strategically choosing the points with respect to the percentiles of simulated processes significantly reduces the number of points needed while keeping the accuracy at a reasonable level.

Next, we need to predict τ_1 based on R_1 , R_0 , τ_0 and so on. The predicted regime $\hat{\tau}$ is the one having the largest probability given the information on prices up to time t , i.e. the *most probable regime* given by (4). Note that this methodology introduces a bias. We also studied the less biased approach of weighting the regimes proportionally to η_t , but since the results were comparable and did not lead to any significant improvement, they are omitted from the analysis. For more details on regime predictions, see Section 2.1. Then, according to (7) and (16), the optimal hedging weights φ_t for period $[t-1, t)$, for $k = n, \dots, 1$, are approximated by

$$\hat{\varphi}_t = \alpha_t(\check{S}_{t-1}, Y_{t-1}, \hat{\tau}_{t-1}) - D^{-1}(\check{S}_{t-1})\check{V}_{t-1}h_t(Y_{t-1}, \hat{\tau}_{t-1}). \quad (18)$$

V_0 is approximated by $C_0(S_0, \hat{\tau}_0, 0)$ while the remaining monies, $V_0 - \hat{\varphi}_1^\top S_0$, are invested in the riskless asset. Next, as S_1 is observed, one first computes the actual portfolio value V_1 , then predicts the current regime τ_1 and finally approximates the optimal weights φ_2 . This process is iterated until expiration of the option.

3.1.1. Using regime predictions

Here, we obtain option prices and strategies that depend on the unobservable regimes τ , since (S_t, τ_t) is a Markov chain. However, [25] proposed a very interesting approach: they showed that (S_t, η_t) is Markov, so one can obtain prices and hedging strategies depending on (S_t, η_t) instead. This makes sense financially. However, this new Markov chain lives in a $l + d - 1$ -dimensional space, because the values of η_t belong to the simplex $\mathcal{S}_l = \{(x_1, \dots, x_l); \ x_i \geq 0, x_1 + \dots + x_l = 1\}$. [25] considered only 2 regimes and one asset, so the real dimension is 2. When $l > 3$, this becomes numerically intractable.

3.2. Optimal hedging vs delta-hedging

In optimal hedging, the strategy φ takes into account the hedging error at each time period, as exemplified by equation (7), since it depends explicitly on the previous value of the portfolio. For delta-hedging, this is not the case since it is assumed that the hedging error is 0 at each period, so there is no correction term depending on the value of the portfolio.

3.3. Simulated hedging errors

To assess the proposed strategy's accuracy, we simulated 10000 trajectories under ARHMM and we hedged identical options under different hedging strategies. To be realistic, the parameters were taken from Table 3. The hedging methodologies are the classical Black-Scholes delta-hedging (B&S) and optimal hedging under ARHMM (OH-ARHMM), HMM (OH-HMM) and Gaussian (OH-B&S) returns (i.e., considering only 1 regime). We also compared semi-exact approximation to Monte-Carlo. The option in question is a call with S_0 and K equal to 100, risk-free rate $r = 0.01$, 3 month maturity (63 days) with daily hedging. Hedging error statistics are given in Table 8, while the estimated densities are displayed in Figure 8. OH-ARHMM achieves a 33% reduction in RMSE compared to B&S and OH-B&S and a 26% to OH-HMM. The latter is quite impressive, as it highlights how big of an impact the autoregressive dynamic has.

4. Out-of-sample vanilla pricing and hedging

4.1. Methodology

To exhibit the behavior of the hedging protocols, we buy and sell vanilla options on the S&P 500 depending on how the market prices compare with

Table 8. Statistics for the hedging errors in an autoregressive hidden Markov model, using 10000 simulated portfolios.

	B&S	OH-B&S	HMM	HMM MC	ARHMM	ARHMM MC
Average	-0.105	-0.084	0.004	0.003	0.025	0.030
Median	-0.236	-0.202	-0.085	-0.086	-0.019	-0.019
Volatility	0.611	0.626	0.559	0.559	0.411	0.412
Skewness	1.715	1.948	1.639	1.629	4.644	4.738
Kurtosis	7.737	9.558	9.464	9.356	71.834	78.219
Minimum	-1.658	-1.658	-4.649	-4.633	-2.477	-2.413
VaR (1%)	-1.110	-1.118	-1.087	-1.086	-0.749	-0.741
VaR (99%)	2.069	2.227	1.987	1.982	1.526	1.531
Maximum	4.886	6.725	6.538	6.520	12.958	14.266
RMSE	0.620	0.632	0.559	0.559	0.411	0.413

Fig. 8. Estimated densities for the hedging errors in an autoregressive hidden Markov model, using 50000 portfolios. Only the semi-exact densities are shown, as they were indiscernible from the Monte Carlo ones.

our theoretical prices. Then, we hedge the positions until expiration. We then assess the impact of model specification on the delta-hedging strategy by examining the statistical properties of the hedging error and of the strategy's returns. All hedging portfolios are re-balanced on a daily basis, as is often assumed in the volatility timing literature; see e.g., [26]. The market price of an option is defined as the last (i.e. as 4:15 PM EST) midpoint between the bid and the ask. The price of the underlying is its listed close value. For simplicity, we neglect issues related to time-varying discount rates by assuming constant continually compounded daily rates. Risk-free rates, r , are linearly interpolated for a given maturity, n , from the zero-coupon U.S. yield curve.

Remark 4.1. For the implementation, we chose to present only the results using the most probable regime for the computation of the hedging strategy. These results are a little bit better than those obtained by weighting the hedging strategy according to $\eta_t(1), \dots, \eta_t(l)$ at period t .

4.1.1. The underlying asset

We make the reasonable assumption that the spot S&P 500 is investable and tradable at a minimal cost. The forward rate is retrieved for the maturities of interest directly from the option data at hand, as proposed by [27]. From put-call parity, the option implied forward value at n , F_n , is

$$F_n = (\tilde{C}(\tilde{K}, n) - \tilde{P}(\tilde{K}, n))e^{r_n n} + \tilde{K},$$

where $C(\tilde{K}, T)$ and $P(\tilde{K}, T)$ are respectively the call and put *market* values expiring at T with strike K and $\tilde{K}$ is the at-the-money strike value minimizing $|C(\tilde{K}, T) - P(\tilde{K}, T)|$ for all strikes offered by the exchange. We use at-the-money options because they are the most liquid and are thus less likely to provide *cash-and-carry* type arbitrage opportunities. We then compute the daily forward rate as $f_n = \frac{1}{n} \log(F_n/S_0)$ and the associated daily discounting factor $\beta = e^{-f_n}$, which reflects the current risk-free return on capital net of the implied continuous dividend yield.

4.1.2. Option dataset

Exchange-traded options on the S&P 500 are European, heavily traded and have a high number of strikes and maturities. To assess the accuracy of our model, we will analyze two periods with very different characteristics: the 2008 Financial Crisis, and a chunk of the recent recovery. Dates range from 09/24/2007 to 09/20/2009 and from 09/23/2013 to 07/08/2015, respectively. This will help us discern the impact on hedging and pricing when a dramatic regime change occurs, in the former, and when it does not, in the latter. In order to minimize the effect of varying maturities, we built the dataset of options having a maturity of about 1 year, more precisely from 231 to 273 trading days till expiration. Also, because *in-the-money* and *out-of-the-money* are less liquid, we only included options where moneyness (strike value divided by the underlying value) is between 0.9 and 1.1. So we were left with a total of 180 options for the first period, and 478 options for the second period. Note that at a given date, more than one option can meet these criteria.

4.1.3. Backtesting

We applied the AR(1) regime-switching optimal hedging methodology with 3 regimes (ARHMM). We chose 3 regimes because it is the number of regimes that seemed the best given the time windows studied, which we will describe in the next paragraph. We compared it to the case with 1 regime and Φ fixed at 0, corresponding to the optimal hedging under the B&S model (OH-B&S).

For each option in the dataset, we estimated the ARHMM parameters on the S&P 500 log-returns with a 500 and 2000 day trailing window. We chose to backtest the methods using 2 estimation windows in order to have a more in-depth understanding of model specifications on pricing and hedging. The 2000 day trailing window always included the previous financial meltdown, i.e., dot-com bubble for our first analysis, and the 2008 financial crisis for the second one. The 500 day trailing window didn't. Similarly, we applied this methodology to all the hedging protocols included in the analysis, which will be introduced below.

From [28], for a given *moneyness*, the value of an option is homogeneous of degree one with respect to the underlying value. Thus, for each inception date, we normalize the option prices, the strike values and the underlying path at an initial S&P 500 value of 100. Results can thus be aggregated through time and interpreted as a percentage of S&P 500. Note that for each inception date, the hedging protocols are applied out-of-sample until maturity. To ensure comparability, OH-B&S assumes the stationary distribution of the ARHMM when the autoregressive parameter $\Phi = 0$. The OH-B&S optimal hedging exposure is derived from an algorithm similar to the one presented in Section 3. Optimal hedging under unconditional distributions is presented in [18]. Both strategies minimize the expected quadratic hedging error under their respective null hypothesis, namely that the returns follow an autoregressive regime-switching model (ARHMM), and a Gaussian model (OH-B&S).

OH-B&S methodology is not to be confused with the classical Black-Scholes delta hedging protocol. Indeed, the terminology only reflects the fact that we hedge and price under the Black-Scholes framework hypothesis, namely that assets follow geometric Brownian motions. Even though the OH-B&S prices converge to the usual Black-Scholes prices as the number of hedging periods tends to infinity, the discrete time hedging strategies will not necessarily be the same. For this reason, the classical Black-Scholes delta-hedging methodology (B&S) is also considered. Similarly to OH-B&S, the B&S volatility is calibrated to the stationary volatility of ARHMM.

We added a final benchmark to our analysis, one that reflects how well the market would hedge the same options, namely the delta-hedging methodology where the volatility is calibrated to the implied volatility at each hedging period (B&S-M). It informs us how well the models compare to market's intuition. The effect of using implied volatility was discussed in [29]. However, his theoretical analysis cannot be performed here.

To recap, we buy and sell options depending on their market value compared to the theoretical prices, and hedge the positions until maturity. We analyze the P&L of the different methodologies, as well as the hedging errors. Two periods are studied: the 2008 Financial Crisis and a chunk of the following recovery spanning from mid-2013 to mid-2015.

4.2. Empirical results

We define the hedging error as the present value of the liability $\beta_n C$ minus the present value of terminal portfolio $\beta_n V_n$. The options' maturity being set to one year, the annualized root-mean-squared hedging error can be computed by $\sqrt{\hat{E}(\beta_n V_n - \beta_n C)^2}$. This realized risk is the empirical counterpart of the quantity we minimized and as such, is the most relevant metric for comparing the different models. Keep in mind that there is a lot of overlap in our dataset, so the hedging error values are not independent, nor identically distributed since the moneyness or other parameters are not constant. Despite these inconveniences, the hedging errors are still useful to compare the models.

Concerning the trading strategy, if the market is overvalued with respect to the model, we sell the option and hedge our position. Thus, the present value of the return is $(C_0 - V_0) - (\beta_n C_n - \beta_n V_n)$. If the market is undervalued, we buy the option and hedge our position. The return will be the negative of the former.

4.2.1. 2008–2009 Financial Crisis

In this section, we will focus on options with inception dates from September 24th 2007 to September 20th 2009. This period is really interesting. In the first part, the market experienced a huge increase in volatility and decrease in returns. In the second part, it is the opposite.

We will first turn our attention to the 500 trailing window case. Table 9 and Figure 9 present the hedging error's statistics and density approximation, while the results of the trading strategy, i.e., the cumulative value

of a portfolio that traded the 90 options, are displayed in Figure 10. The horizontal axis is the cumulative number of options traded in chronological order. In this case, ARHMM is by far the superior methodology. It achieved the best hedging error considering all the metrics for both calls and puts. Furthermore, it is the best trading strategy for both type of options, even though the hedging errors are almost entirely negative in the calls case. Note that the “Bias” statistic refers to the difference between the market price and the theoretical price. Therefore, it is always 0 for the BS-M, since the implied volatility is used.

Note that according to Figure 4, around 2009, the asset spends a lot of time in regime 3, which is characterized by a very large volatility and a very large negative mean; see, e.g., Table 5. When volatility increases and returns turn negative, the puts’ value increase and need to be hedged accordingly. B&S and B&S-M failed to do so, resulting in huge hedging errors and large negative P&L values. This explains the difference between put and calls in Table 9.

Table 9. Hedging error statistics for the 90 calls and the 90 puts traded in the 2008-2009 Financial Crisis with 500 days trailing estimation window.

	Calls				Puts			
	B&S-M	B&S	OH-B&S	ARHMM	B&S-M	B&S	OH-B&S	ARHMM
RMSE	3.87	5.27	4.53	0.61	39.95	42.52	4.53	0.98
Bias	0	-4.52	-4.37	-5.35	0	-1.05	-0.91	-1.87
VaR 1%	-7.64	-12.02	-12.47	-3.17	-28.39	-33.24	-12.47	-3.18
Median	2.9	3.82	2.68	-1.65e-04	30.78	29.77	2.68	0.01
VaR 99%	9.16	8.61	7.57	-1.97e-09	72.9	77.22	7.57	4.75

Similar results are presented in Table 10 and Figures 11 and 12, although the trailing estimation window, previously set to 500 days, is now 2000 days. This estimation window includes another financial crisis, the Dot-com Bubble. The same conclusions as in the previous experience can be drawn.

Table 10. Hedging error statistics for the 90 calls and the 90 puts traded in the 2008-2009 Financial Crisis with 2000 days trailing estimation window.

	Calls				Puts			
	B&S-M	B&S	OH-B&S	ARHMM	B&S-M	B&S	OH-B&S	ARHMM
RMSE	3.87	4.26	3.15	0.33	39.95	40.87	3.15	1.25
Bias	0	-4.86	-4.69	-4.91	0	-1.4	-1.23	-1.44
VaR 1%	-7.64	-4.76	-4.27	-1.43	-28.39	-28.68	-4.27	-1.4
Median	2.9	3.26	1.93	0.01	30.78	27.74	1.93	0.35
VaR 99%	9.16	8.93	7.62	0.33	72.9	74.76	7.62	4.83

Fig. 9. Hedging error density approximation for the 90 calls (a) and 90 puts (b) traded in the 2008-2009 Financial Crisis with 500 days trailing estimation window.

4.2.2. 2013–2015 Bull markets

Our second and last analysis focuses on a part of the recent recovery spanning from September 23th 2013 to August 7th 2015. This period is quite the opposite of a financial crash, with the exception of periods spent in regime 3, which has a very large negative mean and a large volatility. This also affects negatively the performance of hedging error statistics of the B&S methods for the puts. Again, we start with the small estimation window. The results are presented in Table 11 and Figures 13 and 14. Considering the hedging errors, OH-B&S and ARHMM achieved the best and pretty

Fig. 10. Profit & Loss of trading strategy for the 90 calls (a) and 90 puts (b) traded in the 2008-2009 Financial Crisis with 500 days trailing estimation window.

similar statistics for both put and calls. Similarly to the previous experience in Section 4.2.1, B&S and B&S-M replicated poorly the put options.

Table 11. Hedging error statistics for the 239 calls and the 239 puts traded in the 2013-2015 Bull markets with 500 days trailing estimation window.

	Calls				Puts			
	B&S-M	B&S	OH-B&S	ARHMM	B&S-M	B&S	OH-B&S	ARHMM
RMSE	1.09	1.64	0.84	0.99	18.12	11.6	0.84	0.99
Bias	0	0.12	0.22	0.59	0	-4.08	-3.98	-3.62
VaR 1%	-2.63	-2.53	-1.45	-2.47	-41.63	-26.59	-1.45	-2.46
Median	-0.2	0.14	-0.02	-0.01	-12.12	-8.6	-0.02	-0.01
VaR 99%	1.39	4.42	3.7	3.54	9.42	6.78	3.7	3.53

Finally, the results for the longer estimation window case are presented in Table 12 and Figures 15 and 16. This is probably the worst environment for the ARHMM, as the estimation window includes a financial crisis (i.e. 2008-2009 Financial Crisis) and the out-of-sample returns are slow and steady. Because our trading strategy takes into account the actual hedging error according to (7), the simpler models should perform better. In spite of that, ARHMM managed to outperform B&S and B&-M for the put options. The fact that pricing bias for the calls are strongly positive is noteworthy. In theory, the pricing bias should be negative, to account for the risk premium. In this case, it seems that the market was pretty confident about returns and volatility staying low. In insight, it was right.

Lastly, we aggregated the P&L over all the experiences for B&S, OH-B&S and ARHMM in Table 13. For a fair comparison, we normalized the number of traded options in each cases to 100. Remember that the option prices, strike prices and underlying paths are also normalized at an initial S&P 500 value of 100. Impressively, ARHMM accomplished a 106% increase in P&L compared to the second best, OH-B&S, for the 2-year

Fig. 11. Hedging error density approximation for the 90 calls (a) and 90 puts (b) traded in the 2008-2009 Financial Crisis with 2000 days trailing estimation window.

Table 12. Hedging error statistics for the 239 calls and the 239 puts traded in the 2013-2015 Bull markets with 2000 days trailing estimation window.

	Calls				Puts			
	B&S-M	B&S	OH-B&S	ARHMM	B&S-M	B&S	OH-B&S	ARHMM
RMSE	1.09	5.32	4.84	8.57	18.12	14.75	4.84	8.6
Bias	0	4.71	4.54	4.18	0	0.52	0.34	-0.01
VaR 1%	-2.63	-10	-8.93	-31.27	-41.63	-27.66	-8.93	-31.39
Median	-0.2	-4.67	-4.09	-2.82	-12.12	-12.5	-4.09	-2.81
VaR 99%	1.39	-1.79	-1.76	-0.04	9.42	1.41	-1.76	-0.04

trailing window, and is only 9% behind the first for the 8-year case, which is again OH-B&S.

Fig. 12. Profit & Loss of trading strategy for the 90 calls (a) and 90 puts (b) traded in the 2008-2009 Financial Crisis with 2000 days trailing estimation window.

Table 13. Total normalized P&L

Trailing window (years)	B&S	OH-B&S	ARHMM
2	-1286.78	685.32	1409.54
8	-1761.77	594.56	546.38

Overall, by achieving the best Root Mean Square Error (RMSE) two times out of four for both the 2-year and 8-year window, and by being the most profitable strategy three times out of four for the 2-year window and two times out of four for the 8-year window, the ARHMM is the superior hedging protocol. However, practitioners should keep in mind that if the ARHMM is estimated on a window including a financial crisis, they should expect higher hedging errors than the simpler models if returns stay slow and steady. From our results, we strongly suggest using a 2-year trailing window as it consistently achieved an RMSE lower than 1, i.e., the ARHMM can accurately hedge options in a financial crisis.

5. Conclusion

In this paper, we proposed an autoregressive hidden Markov model to fit financial data, and we showed how to implement an optimal hedging strategy when the underlying asset returns follow an autoregressive regime-switching random walk. First, we presented estimation and filtering procedures for the ARHMM. In order to determine the optimal number of regimes, we proposed a novel goodness-of-fit test for univariate and multivariate ARHMM based on the work of [30], [20] and [12].

To illustrate the proposed strategy, we modeled three daily return series of the S&P 500. Using likelihood test, we show that the ARHMM

Fig. 13. Hedging error density approximation for the 239 calls (a) and 239 puts (b) traded in the 2013-2015 Bull markets with 500 days trailing estimation window.

is a much better fit than the classical HMM, particularly because it has the capacity to model mean-reversion. Moreover, we presented the implementation of the discrete-time optimal hedging algorithm minimizing the mean-squared hedging error. Because it performs pricing, we implemented a trading strategy consisting of selling overpriced and buying underpriced options and hedging the position till maturity. Out of eight cases and compared to three other hedging protocols, our strategy achieves the best root-mean-squared hedging error four times and is the most profitable strategy five times. Furthermore, it realized the best total P&L. Because of its

Fig. 14. Profit & Loss of trading strategy for the 239 calls (a) and 239 puts (b) traded in the 2013-2015 Bull markets with 500 days trailing estimation window.

Fig. 15. Hedging error density approximation for the 239 calls (a) and 239 puts (b) traded in the 2013-2015 Bull markets with 2000 days trailing estimation window.

Fig. 16. Profit & Loss of trading strategy for the 239 calls (a) and 239 puts (b) traded in the 2013-2015 Bull markets with 2000 days trailing estimation window.

ability to model regime switches as well as mean-reversion, it would be interesting to see this model applied to multivariate time series. The hedging algorithm can also be applied to multivariate or American options.

Appendix A. Extension of Baum-Welch algorithm

For $i \in \{1, \dots, l\}$ and $t \in \{1, \dots, n\}$, define

$$\lambda_t(i) = P(\tau_t = i | Y_1, \dots, Y_n).$$

Also, for $i, j \in \{1, \dots, l\}$ and $t \in \{1, \dots, n-1\}$, define

$$\Lambda_t(i, j) = P(\tau_t = i, \tau_{t+1} = j | Y_1, \dots, Y_n),$$

and let $\bar{\eta}_t(i)$ be the conditional density of $(Y_{t+1}, \dots, Y_n)$, given Y_t and $\tau_t = i$. Further set $\bar{\eta}_n \equiv 1$. Note that $\lambda_n(i) = \eta_n(i)$, and $\Lambda_n(i, j) = \lambda_n(i)Q_{ij}$, for any $i, j \in \{1, \dots, l\}$. The proof of the following proposition is given in [24].

Proposition A.1. *For all $i, j \in \{1, \dots, l\}$, and any $t \in \{0, \dots, n-1\}$,*

$$\eta_{t+1}(i) = \frac{f_i(Y_{t+1} | Y_t) \sum_{\beta=1}^l \eta_t(\beta) Q_{\beta i}}{\sum_{\alpha=1}^l \sum_{\beta=1}^l f_\alpha(Y_{t+1} | Y_t) \eta_t(\beta) Q_{\beta \alpha}}, \quad (\text{A.1})$$

$$\bar{\eta}_t(i) = \sum_{\beta=1}^l Q_{i\beta} \bar{\eta}_{t+1}(\beta) f_\beta(Y_{t+1} | Y_t), \quad (\text{A.2})$$

$$\lambda_t(i) = \frac{\eta_t(i) \bar{\eta}_t(i)}{\sum_{\alpha=1}^l \eta_t(\alpha) \bar{\eta}_t(\alpha)}, \quad (\text{A.3})$$

$$\Lambda_t(i, j) = \frac{\eta_t(i) Q_{ij} \bar{\eta}_{t+1}(j) f_j(Y_{t+1} | Y_t)}{\sum_{\alpha=1}^l \eta_t(\alpha) \bar{\eta}_t(\alpha)}. \quad (\text{A.4})$$

In particular, $\sum_{\beta=1}^l \Lambda_t(i, \beta) = \lambda_t(i)$, for all $t \in \{0, \dots, n\}$.

Appendix A.1. Estimation of regime-switching models

To describe the EM algorithm for the estimation, suppose that at step $k \geq 0$, one has the parameters $Q, \mu_i, \Phi_i, A_i, i \in \{1, \dots, l\}$. Let $w_t(i) = \lambda_t(i) / \sum_{k=1}^n \lambda_k(i)$, and set $\bar{y}_i = \sum_{t=1}^n w_t(i) y_t$ and $\underline{y}_i = \sum_{t=1}^n w_t(i) y_{t-1}$, $i \in \{1, \dots, l\}$, where λ_t and Λ_t are given in Proposition A.1. Then, at step $k+1$, for $i, j \in \{1, \dots, l\}$, one has

$$Q_{ij}^{(k+1)} = \frac{\sum_{t=1}^n \Lambda_{t-1}(i, j)}{\sum_{\beta=1}^l \sum_{t=1}^n \Lambda_{t-1}(i, \beta)} = \frac{\sum_{t=1}^n \Lambda_{t-1}(i, j)}{\sum_{t=1}^n \lambda_{t-1}(i)}, \quad (\text{A.5})$$

$$\mu_i^{(k+1)} = \left(I - \Phi_i^{(k+1)} \right)^{-1} \left(\bar{y}_i - \Phi_i^{(k+1)} \underline{y}_i \right), \quad (\text{A.6})$$

$$\Phi_i^{(k+1)} = \left\{ \sum_{t=1}^n w_t(i) \left(y_{t-1} - \underline{y}_i \right) \left(y_{t-1} - \underline{y}_i \right)^\top \right\}^{-1} \quad (\text{A.7})$$

$$\times \left\{ \sum_{t=1}^n w_t(i) (y_t - \bar{y}_i) \left(y_{t-1} - \underline{y}_i \right)^\top \right\},$$

$$A_i^{(k+1)} = \sum_{t=1}^n w_t(i) e_{ti} e_{ti}^\top, \quad (\text{A.8})$$

where $e_{ti} = y_t - \bar{y}_i - \Phi_i^{(k+1)} \left(y_{t-1} - \underline{y}_i \right)$, $t \in \{1, \dots, n\}$.

For a proof of these formulas, see [24].

Appendix B. Goodness-of-fit test for ARHMM

Here, we state the main formulas needed to implement the goodness-of-fit test, which can be performed to assess the suitability of a Gaussian AR(1) regime-switching models as well as to select the optimal number of regimes, l^* . The proposed test, based on the work of [31], [20] and [32], uses the Rosenblatt's transform. For conciseness, we detail the implementation for two dimensional Gaussian AR(1) regime-switching models, but the approach can be easily generalized.

Appendix B.1. Rosenblatt's transform

Let $i \in \{1, \dots, l\}$ be fixed, and let R_i be a random vector with density f_i . For any $q \in \{1, \dots, d\}$, denote by $f_{i,1;q}$ the density of $(R_i^{(1)}, \dots, R_i^{(q)})$, and by $f_{i,q}$ the density of $R_i^{(q)}$ given $(R_i^{(1)}, \dots, R_i^{(q-1)})$. Further denote by $F_{i,q}$ the distribution function associated with density $f_{i,q}$. By convention, $f_{i,1}$ denotes the *unconditional* density of $R_i^{(1)}$. Then, the Rosenblatt's transform $x \mapsto T_i(x) = (F_{i,1}(x^{(1)}), F_{i,2}(x^{(1)}, x^{(2)}), \dots, F_{i,d}(x^{(1)}, \dots, x^{(d)}))^\top$ is such that $T_i(R_i)$ is uniformly distributed in $[0, 1]^d$. For example, if f_i is the density of a bivariate Gaussian distribution with mean μ_i and covariance

$$\Sigma_i = \begin{pmatrix} v_i^{(1)} & \rho_i \sqrt{v_i^{(1)} v_i^{(2)}} \\ \rho_i \sqrt{v_i^{(1)} v_i^{(2)}} & v_i^{(2)} \end{pmatrix},$$

$f_{i,2}$ is the density of a Gaussian distribution with mean $\mu_i^{(2)} + \beta_i (y_i^{(1)} - \mu_i^{(1)})$ and variance $v_i^{(2)} (1 - \rho_i^2)$, with $\beta_i = \rho_i \sqrt{v_i^{(2)} / v_i^{(1)}}$. These results can easily be extended to the Gaussian AR(1) distribution.

However, for regime-switching random walks models, past returns must also be included in the conditioning information set. For any $x^{(1)}, \dots, x^{(d)} \in \mathbb{R}$, the (d -dimensional) Rosenblatt's transform Ψ_t corresponding to the density (5) conditional on $x_1, \dots, x_{t-1} \in \mathbb{R}^d$ is given by

$$\Psi_t^{(1)}(x_t^{(1)}) = \Psi_t^{(1)} \left(x_1, \dots, x_{t-1}, x_t^{(1)} \right) = \sum_{i=1}^l W_{t-1}(i) F_{i,1}(x_t^{(1)})$$

and

$$\begin{aligned} \Psi_t^{(q)} \left(x_1^{(1)}, \dots, x_t^{(q)} \right) &= \Psi_t^{(q)} \left(x_1, \dots, x_{t-1}, x_t^{(1)}, \dots, x_t^{(q)} \right) \\ &= \frac{\sum_{i=1}^l W_{t-1}(i) f_{i,1:q-1} \left(x_t^{(1)}, \dots, x_t^{(q-1)} \right) F_{i,q}(x_t^{(q)})}{\sum_{i=1}^l W_{t-1}(i) f_{i,1:q-1} \left(x_t^{(1)}, \dots, x_t^{(q-1)} \right)} \end{aligned}$$

for $q \in \{2, \dots, d\}$. Suppose $R_1, \dots, R_n$ is a sample of size n of d -dimensional vectors drawn from a joint (continuous) distribution P . Also, let $\mathcal{P}$ be the parametric family of Gaussian AR(1) regime-switching models with l regimes. Formally, the hypothesis to be tested is

$$\mathcal{H}_0 : P \in \mathcal{P} = \{P_\theta; \theta \in \Theta\} \quad \text{vs} \quad \mathcal{H}_1 : P \notin \mathcal{P}.$$

Under $\mathcal{H}_0$, it follows that

$$U_1 = \Psi_1(R_1, \theta), U_2 = \Psi_2(R_1, R_2, \theta), \dots, U_n = \Psi(R_1, \dots, R_n, \theta)$$

are independent and uniformly distributed over $[0, 1]^d$, where $\Psi_1(\cdot, \theta), \dots, \Psi_n(\cdot, \theta)$ are the Rosenblatt's transforms conditional on the set of parameters $\theta \in \Theta$. Since θ is unknown, it must be estimated by some θ_n . Then, the *pseudo-observations*, $(\hat{U}_1 = \Psi_1(R_1, \theta_n), \dots, \hat{U}_n = \Psi_n(R_1, \dots, R_n, \theta_n))$ are *approximately* uniformly distributed over $[0, 1]^d$ and *approximately* independent. We propose a test statistic based on these pseudo observations.

Appendix B.2. Test statistic

The test statistic is based on the empirical process $D_n(u) = \frac{1}{n} \sum_{t=1}^n \prod_{q=1}^d \mathbb{I}(\hat{U}_t^{(q)} \leq u^{(q)})$, $u \equiv (u^{(1)}, \dots, u^{(d)}) \in [0, 1]^d$.

To test $\mathcal{H}_0$ against $\mathcal{H}_1$ we propose a Cramér-von Mises type statistic:

$$\begin{aligned} S_n &\equiv B_n(\hat{U}_1, \dots, \hat{U}_n) = n \int_{[0,1]^d} \left\{ D_n(u) - \prod_{q=1}^d u^{(q)} \right\}^2 du \\ &= \frac{1}{n} \sum_{t=1}^n \sum_{k=1}^n \prod_{q=1}^d \left\{ 1 - \max \hat{U}_t^{(q)}, \hat{U}_k^{(q)} \right\} - \frac{1}{2^{d-1}} \sum_{t=1}^n \prod_{q=1}^d (1 - \hat{U}_t^{(q)2}) + \frac{n}{3^d}. \end{aligned}$$

Since $\hat{U}_i$ is almost uniformly distributed under $\mathcal{H}_0$, large values of S_n should lead to rejection of $\mathcal{H}_0$. Unfortunately, the limiting distribution of the test statistic depends on the unknown parameter θ . This is why we use the parametric bootstrap to compute P -values. Its validity has been shown for a wide range of assumptions in [20]. These results were recently extended to dynamic models in [16], including regime-switching random walks.

Appendix B.3. Parametric bootstrap algorithm

- (1) For a given number of regimes, estimate parameters with θ_n computed from the EM algorithm applied to $(R_1, \dots, R_n)$.
- (2) Compute the test statistic $S_n = B_n(\hat{U}, \dots, \hat{U}_n)$, from the estimated pseudo observations, $\hat{U}_i = \Psi_i(R_1, \dots, R_i, \theta_n)$, for $i \in \{1, \dots, n\}$.
- (3) For some large integer N (say 1000), repeat the following steps for every $k \in 1, \dots, N$:
 - (a) Generate a random sample $\{R_1^k, \dots, R_n^k, \theta_n^k\}$ from distribution P_{θ_n} .
 - (b) Compute θ_n^k by applying the EM algorithm to the simulated sample, $R_1^k, \dots, R_n^k$.
 - (c) Let $\hat{U}_i^k = \Psi_i(R_1^k, \dots, R_i^k, \theta_n^k)$ for $i \in 1, \dots, n$, and finally compute $S_n^k = B_n(\hat{U}_1^k, \dots, \hat{U}_n^k)$.

Then, the approximated P -value for the test based on the Cramér von Mises statistic S_n is given by $\frac{1}{N} \sum_{k=1}^N \mathbb{I}(S_n^k > S_n)$.

References

- [1] J. Cox and S. Ross, The valuation of options for alternative stochastic processes, *Journal of Financial Economics* **3**, 145–166 (1976).
- [2] F. Black and M. Scholes, The pricing of options and corporate liabilities, *Journal of Political Economy* **81**, 637–654 (1973).
- [3] G. W. Schwert, Why does stock market volatility change over time?, *Journal of Finance* **44**, 1115–1153 (1989).

- [4] J. D. Hamilton and G. Lin, Stock market volatility and business cycle, *Journal of Applied Econometrics*, 573–593 (1996).
- [5] M. Schweizer, Mean-variance hedging for general claims, *Ann. Appl. Probab.* **2**, 171–179 (1992).
- [6] M. Schweizer, Variance-optimal hedging in discrete time, *Math. Oper. Res.* **20**, 1–32 (1995).
- [7] J. D. Hamilton, A new approach to the economic analysis of non-stationary time series and the business cycle, *Econometrica*, 357–384 (1989).
- [8] J. D. Hamilton, Analysis of time series subject to changes in regime, *J. Econometrics* **45**, 39–70 (1990).
- [9] P. Wu and R. J. Elliott, Parameter estimation for a regime-switching mean-reverting model with jumps, *Int. J. Theor. Appl. Finance* **8**, 791–806 (2005).
- [10] R. J. Elliott, H. Miao and Z. Wu, An asset pricing model with mean reversion and regime switching stochastic volatility, in *The Oxford handbook of nonlinear filtering*, (Oxford Univ. Press, Oxford, 2011) pp. 960–989.
- [11] R. J. Elliott, L. Chan and T. K. Siu, Option valuation under a regime-switching constant elasticity of variance process, *Appl. Math. Comput.* **219**, 4434–4443 (2013).
- [12] B. Rémillard, A. Hocquard, H. Lamarre and N. A. Papageorgiou, Option pricing and hedging for discrete time regime-switching model, *Modern Economy* **8**, 1005–1032 (September 2017).
- [13] S. Shi and A. S. Weigend, Taking time seriously: hidden Markov experts applied to financial engineering, *Computational Intelligence for Financial Engineering (CIFEr)*, 244 (1997).
- [14] A. P. Dempster, N. M. Laird and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. Roy. Statist. Soc. Ser. B* **39**, 1–38 (1977).
- [15] C. Genest, J.-F. Quessy and B. Rémillard, Goodness-of-fit procedures for copula models based on the integral probability transformation, *Scand. J. Statist.* **33**, 337–366 (2006).
- [16] B. Rémillard, *Validity of the Parametric Bootstrap for Goodness-of-Fit Testing in Dynamic Models*, tech. rep., SSRN Working Paper Series No. 1966476 (2011).
- [17] E. F. Fama, The behavior of stock-market prices, *Journal of Business* **38**, 34–105 (1965).
- [18] B. Rémillard, *Statistical Methods for Financial Engineering* Chapman and Hall/CRC Financial Mathematics Series, Chapman and Hall/CRC Financial Mathematics Series (Taylor & Francis, 2013).

- [19] O. Cappé, E. Moulines and T. Rydén, *Inference in Hidden Markov Models* Springer Series in Statistics, Springer Series in Statistics (Springer, New York, 2005).
- [20] C. Genest and B. Rémillard, Validity of the parametric bootstrap for goodness-of-fit testing in semiparametric models, *Ann. Inst. H. Poincaré Sect. B* **44**, 1096–1127 (2008).
- [21] N. Papageorgiou, B. Rémillard and A. Hocquard, Replicating the properties of hedge fund returns, *Journal of Alternative Investments* **11**, 8–38 (2008).
- [22] B. Rémillard and S. Rubenthaler, Optimal hedging in discrete time, *Quantitative Finance* **13**, 819–825 (2013).
- [23] B. Rémillard and S. Rubenthaler, *Option pricing and hedging for regime-switching geometric Brownian motion models*, working paper series, SSRN Working Paper Series No. 2599064 (2016).
- [24] M. Caccia and B. Rémillard, *Option pricing and hedging for discrete time autoregressive hidden Markov model*, Tech. Rep. G-2017-60, GERAD (2017).
- [25] P. François, G. Gauthier and F. Godin, Optimal hedging when the underlying asset follows a regime-switching Markov process, *European Journal of Operational Research* **237**, 312–322 (2014).
- [26] J. Fleming, C. Kirby and B. Ostdiek, The economic value of volatility timing, *The Journal of Finance* **56**, 329–352 (2001).
- [27] A. Buraschi and J. Jackwerth, The price of a smile: hedging and spanning in option markets, *The Review of Financial Studies* **14**, 495–527 (2001).
- [28] R. C. Merton, Theory of rational option pricing, *The Bell Journal of Economics and Management Science* **4**, 141–183 (1973).
- [29] P. Carr, *FAQ's in option pricing theory*, tech. rep., NYU (2002).
- [30] J. Bai, Testing parametric conditional distributions of dynamic models, *The Review of Economics and Statistics* **85**, 531–549 (2003).
- [31] F. X. Diebold, T. A. Gunther and A. S. Tay, Evaluating density forecasts with applications to financial risk management, *International Economic Review* **39**, 863–883 (1998).
- [32] B. Rémillard, Tests of independence, in *International Encyclopedia of Statistical Science*, ed. M. Lovric (Springer Berlin Heidelberg, 2011) pp. 1598–1601.